

UNCERTAINTY IN FEDERATED GRANGER CAUSALITY: FROM ORIGINS TO SYSTEMIC CONSEQUENCES

Anonymous authors

Paper under double-blind review

ABSTRACT

Granger Causality (GC) provides a rigorous framework for learning causal structures from time-series data. Recent federated variants of GC have targeted distributed infrastructure applications (e.g., smart grids) with distributed clients that generate high-dimensional data bound by data-sovereignty constraints. However, Federated GC algorithms only yield deterministic point estimates of causality and neglect uncertainty. This paper establishes the first methodology for rigorously quantifying uncertainty and its propagation within federated GC frameworks. We systematically classify sources of uncertainty, explicitly differentiating aleatoric (data noise) from epistemic (model variability) effects. We derive closed-form recursion expressions modeling the evolution of uncertainty through client-server interactions, and identify four novel cross-covariance components that couple data uncertainties with model parameter uncertainties across the federated architecture. Moreover, we define rigorous convergence conditions for these uncertainty recursions and obtain explicit steady-state variances for both server and client model parameters. More importantly, our convergence analysis demonstrates that steady-state variances depend exclusively on client data statistics, thus eliminating the dependence on initial epistemic priors and enhancing robustness. Empirical evaluations on synthetic benchmarks and real-world industrial datasets demonstrate that explicitly characterizing uncertainty significantly improves the reliability and interpretability of federated causal inference. These results enable robust root-cause analysis in safety-critical privacy-constrained and distributed infrastructures.

1 INTRODUCTION

Complex industrial systems such as smart grids, distributed manufacturing, and supply chain networks are comprised of geographically distributed, tightly coupled system components (clients) that interact dynamically and require seamless coordination for safe and reliable operations. The resulting interactions also mean that any disruptions, such as equipment failures or cyber intrusions, can propagate quickly, triggering cascading effects across the entire system. Mitigating these challenges requires risk-aware causal models that can identify causal influences between clients and quantify their associated uncertainties. This ensures that limited resources for mitigation are efficiently allocated to prevent or minimize cascading disruptions, especially in safety-critical applications.

Granger Causality (GC) is well-suited for learning causal structures from time-series data Shojaie & Fox (2022); Wismüller et al. (2021); Tank et al. (2021). Conventional GC methods rely on data centralization. Clients in our problem setting cannot readily share their data because they are equipped with modern sensors that generate high-dimensional data, and any downsampling or aggregation of this data destroys critical diagnostic information. Additionally, many clients are bound by data-privacy stipulations.

Recent advances in **federated Granger causality** (FedGC) partly address these challenges by exchanging only sufficient statistics rather than raw data Mohanty et al. (2025). Nonetheless, existing FedGC methods provide only deterministic point estimates and neglect the uncertainty introduced by data noise and model variability. To the best of our knowledge, this is the first methodology that rigorously quantifies uncertainty and its propagation within FedGC frameworks.

Without quantifying uncertainties in causal models, decision-makers may underestimate risks, leading to catastrophic failures. For instance, during the 2003 Northeast blackout Muir & Lopatto (2004), control room alarms failed to trigger, and operators believed the system was operating normally. This binary fault detection approach ignored uncertainty and missed early warning signs. Had the variances on inter-utility causal interactions been monitored rather than relying on single-point alarm estimates, operators might have seen a widening of predictive confidence intervals, providing sufficient warning to mitigate the cascade. Unfortunately, the cascade resulted in a massive outage.

This paper establishes the first theoretical foundation for analyzing how uncertainties propagate in the FedGC framework. We address three fundamental questions:

1. **Sources:** What are the primary contributors to uncertainty in federated causal learning?
2. **Propagation:** How do uncertainties compound as clients and servers iteratively optimize loss?
3. **Impact:** What is the cumulative effect of these uncertainties on the fidelity of causal estimates?

Contributions: The key technical contributions of our paper are:

- A systematic classification of uncertainty sources in federated causal learning, distinguishing between *data noise* (aleatoric), and *model variability* (epistemic) effects.
- Closed-form propagation expressions that track how these uncertainty move through *client-to-server*, and *server-to-client* communication channels, and *within the client*, and *within the server* paths, revealing four previously unrecognized *cross-covariance terms* that couple data and parameters across the FedGC framework.
- Spectral-radius based conditions that guarantee convergence of all covariance recursions leading to explicit solutions for the steady-state variances of the server and client models.
- Theoretical results on convergence, showing that these steady-state variances depend only on raw data statistics of the clients (aleatoric), thereby eliminating any influence of the initial epistemic priors in the FedGC framework.

2 LITERATURE REVIEW

In centralized settings, uncertainty propagation under gradient-based optimization has been studied extensively. Recent works Wang et al. (2023); Durasov et al. (2024); Gawlikowski et al. (2023) analyze how noise from stochastic gradients or perturbations propagates through updates, while others Chan et al. (2024); Huseljic et al. (2021); Meinert et al. (2023); Hofman et al. (2024) formalize the coupling between aleatoric and epistemic uncertainties Hüllermeier & Waegeman (2021). These analyses assume centralized training with full data access and do not generalize to federated settings.

In federated learning, uncertainty quantification has advanced mainly through Bayesian methods, including federated Bayesian neural networks Yurochkin et al. (2019), personalized inference Kotelevskii et al. (2022); Zhang et al. (2022), ensemble strategies Chen & Chao (2020), and Monte Carlo dropout Park et al. (2022). While effective for horizontal IID data Yang et al. (2019), these approaches rely on approximate posteriors (variational inference, sampling, dropout) and typically treat aleatoric and epistemic uncertainties independently. VertiBayes Van Daalen et al. (2024) extends Bayesian inference to vertical partitions but remains static, lacking temporal dynamics and client-server uncertainty interactions. Beyond Bayesian approaches, federated Kalman filtering Xing & Xia (2016); Baucas & Spachos (2023) and distributed state estimation Korres (2010); Primadianto & Lu (2016) propagate uncertainty recursively but ignore cross-client causality. Causal federated works such as FedGC Mohanty et al. (2025) focus on deterministic recovery, while variational inference methods Guo et al.; Vedadi et al. (2024) approximate posteriors without modeling cross-covariances. [Recent horizontally federated causal-discovery methods, including Guo et al. \(2024b;a\), recover causal graphs under IID partitions but do not extend to vertically partitioned dynamical systems with temporal coupling.](#) Robust aggregation Pillutla et al. (2022); Li et al. (2021) mitigates heterogeneity but remains agnostic to compounding uncertainty in causal graphs.

Our work departs from these directions by unifying causal structure learning with federated uncertainty propagation. Unlike Bayesian FL methods that assume horizontal partitioning or decouple uncertainties, we explicitly trace how aleatoric noise and epistemic variability interact via cross-covariance terms. This captures uncertainty propagation in federated time-series data while providing a causal interpretation of client-server interactions.

3 PRELIMINARIES

State Space Model. The underlying system is modeled as a linear time-invariant (LTI) state–space:

$$h^t = A h^{t-1} + w^t, \quad y^t = C h^t + v^t \quad (1)$$

where $h^t \in \mathbb{R}^p$ are the **latent low-dimensional states** and $y^t \in \mathbb{R}^d$ are the **measured high-dimensional data** at time t with $d \gg p$. A and C are the constant state-transition and observation matrices; w^t and v^t are zero-mean i.i.d. Gaussian system and measurement noise with covariances Q and R , respectively. We make the following assumptions about the LTI state space model:

1. There are M subsystems such that $h^t = [h_1^t, \dots, h_M^t]$, $y^t = [y_1^t, \dots, y_M^t]$. Each of these subsystems is a client for our problem setting. The states are such that, $h_m^t \in \mathbb{R}^{p_m}$, and $y_m^t \in \mathbb{R}^{d_m}$ with $d_m \gg p_m$, and $\sum_{m=1}^M p_m = p$, and $\sum_{m=1}^M d_m = d$.
2. The observation matrix C is block-diagonal i.e., $C = \text{diag}(C_{11}, \dots, C_{MM})$ where each block $C_{mm} \in \mathbb{R}^{d_m \times p_m}$. The block C_{mm} is known at client m .
3. The state-transition matrix A is not block-diagonal i.e., $\exists n \neq m$ s.t., $A_{mn} \neq 0$ with $A_{mn} \in \mathbb{R}^{p_m \times p_n}$. Each client m knows (or can locally compute) its diagonal block A_{mm} , while the off-diagonal blocks $A_{mn} \forall n \neq m$ are **unknown**.

Granger Causality. A time series h_n is said to *Granger-cause* another series h_m if the inclusion of past values of h_n improves the prediction of h_m . In the state-space setting, this notion is captured by the off-diagonal entries of the state-transition matrix A . For instance, in a system given by,

$$\begin{pmatrix} h_1^t \\ h_2^t \end{pmatrix} = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} h_1^{t-1} \\ h_2^{t-1} \end{pmatrix} + \begin{pmatrix} w_1^t \\ w_2^t \end{pmatrix} \quad (2)$$

the series h_2 Granger-causes h_1 precisely when $A_{12} \neq 0$. More generally, $A_{mn} \neq 0$ indicates that past values of h_n influence the future of h_m , revealing a directed causal link from client n to m . Estimating A_{mn} in a decentralized system is the goal of the federated GC framework.

3.1 FEDERATED GRANGER CAUSALITY

FedGC is a server-client framework with M different clients having **unknown interdependencies** across them. We discuss the details of the client and server models in the FedGC framework.

Client Model. Each client m models its subsystem as an LTI state–space using only the local (diagonal) blocks $A_{mm} \in \mathbb{R}^{p_m \times p_m}$ and $C_{mm} \in \mathbb{R}^{d_m \times p_m}$. A Kalman filter based “*local client model*” is used to compress the high-dimensional data (measurement) $y_t^{(m)} \in \mathbb{R}^{d_m}$ into a low-dimensional state estimate $\hat{h}_t^{(m)} \in \mathbb{R}^{p_m}$ with $d_m \gg p_m \forall m \in \{1, \dots, M\}$ such that:

$$h_{m,c}^t = A_{mm} \hat{h}_{m,c}^{t-1}, \quad \hat{h}_{m,c}^t = h_{m,c}^t + K_m (y_m^t - C_{mm} (h_{m,c}^t)_c) \quad (3)$$

where K_m is the Kalman gain computed from $(A_{mm}, C_{mm}, Q_{mm}, R_{mm})$. The local client model (Kalman Filter) ignores cross-client dynamics A_{mn} ($n \neq m$). In order to compensate for these cross-client dynamics, the client states are *augmented* with machine learning (ML) models to obtain an “*augmented client model*” as follows:

$$\hat{h}_{m,a}^t = \hat{h}_{m,c}^t + \theta_m y_m^t, \quad h_{m,a}^t = A_{mm} \hat{h}_{m,a}^t \quad (4)$$

where $\theta_m \in \mathbb{R}^{p_m \times d_m}$ is a learnable matrix that captures missing cross-client effects. The loss function $(L_m)_a$ is optimized at client m with θ_m being learned using Gradient-descent based update,

$$(L_m)_a = \|y_m^t - C_{mm} h_{m,a}^t\| \text{ s.t.}, \quad \theta_m^{t+1} = \theta_m^t - \eta_1 \nabla_{\theta_m^t} (L_m)_a - \eta_2 \nabla_{\theta_m^t} L_s \quad (5)$$

In the above equation, η_1, η_2 are the learning rates corresponding to client m ’s loss function $(L_m)_a$ and server model’s loss function L_s respectively (the server model is discussed in next paragraph).

Remark 3.1. In the gradient update equation above, while, $\nabla_{\theta_m^t} (L_m)_a$ can be computed locally at the client m , the gradient $\nabla_{\theta_m^t} L_s$ needs information from the server. However, the server cannot directly compute $\nabla_{\theta_m^t} L_s$ as the gradient is w.r.t. client model parameter θ_m . Therefore, the FedGC

framework adopts chain rule to decompose it such that: $\nabla_{\theta_m^t} L_s = (\nabla_{\hat{h}_{m,a}^t} L_s)(\nabla_{\theta_m^t} \hat{h}_{m,a}^t)$. A key advantage of this decomposition is that the first factor i.e., $\nabla_{\hat{h}_{m,a}^t} L_s$ can be communicated from the server to client m , while $\nabla_{\theta_m^t} \hat{h}_{m,a}^t$ can be computed locally at client m .

Server Model. The server collects from each client m the pair $(\hat{h}_{m,c}^{t-1}, \hat{h}_{m,a}^{t-1})$ and stacks them into

$$H_c^{t-1} = [\hat{h}_{1,c}^{t-1}, \dots, \hat{h}_{M,c}^{t-1}], \quad H_a^t = [A_{11}\hat{h}_{1,a}^{t-1}, \dots, A_{MM}\hat{h}_{M,a}^{t-1}] \quad (6)$$

The diagonal matrix $\text{diag}(A_{11}, \dots, A_{MM})$ is assumed to be known at the server. However, the off-diagonal blocks (representing the GC) $A_{mn} \forall n \neq m$ are **unknown**. The server model’s goal is to predict the next-step state H_s^t as follows:

$$H_s^t = [h_{1,s}^t, \dots, h_{M,s}^t] \text{ s.t.}, \quad h_{m,s}^t = A_{mm}\hat{h}_{m,c}^{t-1} + \sum_{n \neq m} \hat{A}_{mn} \hat{h}_{n,c}^{t-1}. \quad (7)$$

where the estimated causality $\hat{A}_{mn} \forall n \neq m$ are learned by minimizing the server loss function L_s with gradient-descent based updating used to learn $\hat{A}_{mn} \forall n \neq m$ with server learning rate γ s.t.

$$L_s = \|H_a^t - H_s^t\|_2^2, \quad \hat{A}_{mn}^{t+1} = \hat{A}_{mn}^t - \gamma \nabla_{\hat{A}_{mn}^t} L_s \quad (8)$$

The server then sends $\nabla_{\hat{h}_{m,a}^t} L_s$ to client m for subsequent client parameter updating.

4 PROBLEM FORMULATION

The reliance of the FedGC framework on deterministic point estimates of the Granger causal parameters \hat{A}_{mn} ($\forall n \neq m$) and the client parameters θ_m overlooks the influence of data noise and model variability. In this work, we rigorously characterize and quantify how uncertainty propagates through the federated pipeline. Our analysis is anchored around three core objectives:

1. Characterize the **sources of uncertainty**—namely, aleatoric noise from client data and epistemic variability in model parameters.
2. Derive exact **propagation recursions** that track how these uncertainties evolve across client–server communication and local/global updates within the client/server.
3. Analyze the steady-state **impact of uncertainty**, proving that under mild assumptions, the resulting variances depend only on data noise and fixed gains, not on initial model priors.

5 SOURCES OF UNCERTAINTY

In this section, we partition the stochastic elements of FedGC into two disjoint categories: **(i)** *aleatoric* noise originating in the data (measurements) collected by each client, and **(ii)** *epistemic* variance arising from incomplete knowledge of both client-side and server-side model parameters.

(1) Aleatoric. Aleatoric uncertainty captures the irreducible measurement noise ϵ_m^t in client m ’s data i.e., y_m^t . When a client updates its parameter θ_m^t via gradient descent, this data noise ϵ_m^t propagates directly into the gradient $\nabla_{\theta_m^t} (L_m)_a$ and hence into the variance of θ_m^t . Likewise, it also enters the augmented state $\hat{h}_{m,a}^t = \hat{h}_{m,c}^t + \theta_m^t y_m^t$. Since $\hat{h}_{m,a}^t$ is communicated to the server from clients, this data noise also influences the estimation of Granger causality $\hat{A}_{mn}^t \forall n \neq m$ (also called the *server model parameter*). Furthermore, this data noise affects the server loss L_s and shows up in the gradients $\nabla_{\hat{h}_{m,a}^t} L_s \forall m$ communicated from the server to the clients.

(2) Epistemic. Epistemic uncertainty reflects our lack of knowledge about the model parameters—both the client parameters θ_m and the server parameters A_{mn} . We assume both of these parameters to be **random variables** in our problem setting. Sampling from a prior $\theta_m^0 \sim \mathcal{D}_1(\mu_{\theta_m}^0, \Sigma_{\theta_m}^0)$ and $A_{mn}^0 \sim \mathcal{D}_2(\mu_{A_{mn}}^0, \Sigma_{A_{mn}}^0)$ where, \mathcal{D}_1 and \mathcal{D}_2 are any location-scale distributions, we refine θ_m^0 , and \hat{A}_{mn}^0 using gradient-descent updates. By accumulating sufficient gradient-descent iterations, we reduce this epistemic uncertainty and thereby increase our confidence in the estimated causality.

Table 1: Summary of notation (client index m , time index t)

Symbol	Meaning	Shape / Statistics
y_m^t	Raw data for client m at time t	$\in \mathbb{R}^{d_m}$; $\mu_{y_m}^t = \mathbb{E}[y_m^t]$, $\Sigma_{y_m}^t = \text{Var}(y_m^t)$
θ_m^t	Model parameter at client m	$\in \mathbb{R}^{p_m \times d_m}$
v_m^t	Vectorised θ_m^t i.e., $v_m^t = \text{Vec}(\theta_m^t)$	$\mu_{\theta_m}^t = \mathbb{E}[v_m^t]$, $\Sigma_{\theta_m}^t = \text{Var}(v_m^t)$
Ω_m^t	Parameter-data covariance at client m	$\text{Cov}(v_m^t, y_m^t)$
Λ_m^t	Client parameter-state covariance	$\text{Cov}(v_m^t, \hat{h}_{m,a}^t)$
Ψ_{mn}^t	Server-client parameter covariance	$\text{Cov}(a_{mn}^t, v_m^t)$
Γ_{mn}^t	Cross-covariance between a_{mn}^t and $\hat{h}_{m,a}^t$	$\text{Cov}(a_{mn}^t, \hat{h}_{m,a}^t)$
$\hat{h}_{m,a}^t$	Augmented state estimate at client m	$\Sigma_{h_m}^t = \text{Var}(\hat{h}_{m,a}^t)$
\hat{A}_{mn}^t	Server parameter estimate ($n \rightarrow m$)	$\in \mathbb{R}^{p_m \times p_n}$, $\forall n \neq m$
a_{mn}^t	Vectorised \hat{A}_{mn}^t i.e., $a_{mn}^t = \text{Vec}(\hat{A}_{mn}^t)$	$\Sigma_{A_{mn}}^t = \text{Var}(a_{mn}^t)$
L_s	Loss function at server	$\in \mathbb{R}^1$ (Scalar)
$(L_m)_a$	Loss function at client m	$\in \mathbb{R}^1$ (Scalar)
$g_{m,s}^t$	Gradient of server loss L_s w.r.t $\hat{h}_{m,a}^t$	$\in \mathbb{R}^{p_m}$; $\text{Var}(g_{m,s}^t)$
γ	Learning rate of the server model	$\in \mathbb{R}^1$ (Scalar)
η_1, η_2	Learning rates of the client model	$\in \mathbb{R}^1$ (Scalars)

Assumptions. Based on the sources of uncertainty, we make the following assumptions:

- **(A1) Stochasticity.** For every client m , the *client model parameter* θ_m^t , and *client data* y_m^t are random variables. **The local client state $\hat{h}_{m,c}^t$ is deterministic.** Consequently, the only randomness entering the augmented states $\hat{h}_{m,a}^t = \hat{h}_{m,c}^t + \theta_m^t y_m^t$ comes from θ_m^t and y_m^t . In the server model, the Granger causal estimation $\hat{A}_{mn}^t \forall n \neq m$ (also called the *server model parameter*) is random.
- **(A2) Model Parameters.** The server parameters A_{mn}^t , $n \neq m$, are mutually independent across block-rows and times. As a consequence, for any distinct clients $m \neq n$, the induced parameters θ_m^t and θ_n^t have zero cross-covariance at every time t , i.e. $\text{Cov}(\theta_m^t, \theta_n^t) = 0$, $\forall m \neq n$, $\forall t$. We formally establish this result in Lemma D.1, and Proposition D.2.
- **(A3) Prior.** The initial server and client parameters are independent, i.e., $\hat{A}_{mn}^0 \perp\!\!\!\perp \theta_m^0$, and the initial client parameters are uncorrelated, i.e., $\text{Cov}(\theta_m^0, \theta_n^0) = 0 \forall n \neq m$.
- **(A4) Stationarity.** Client data are weakly stationary with time-invariant first and second moments: $\mathbb{E}[y_m^t] = \mu_{y_m}$, $\text{Var}(y_m^t) = \Sigma_{y_m}$, $\forall t$
- **(A5) Noise.** The underlying states h^t in the state–space model (see Eq 1) are noiseless. All stochasticity enters through additive data (measurement) noise: $y_m^t = \mu_{y_m} + \varepsilon_m^t$, with $\mathbb{E}[\varepsilon_m^t] = 0$, and $\text{Var}(\varepsilon_m^t) = \Sigma_{y_m}$. We formally show this using the Proposition D.3.

6 UNCERTAINTY PROPAGATION

Notation. Table 1 gives an overview of the symbols discussed in Sections 3, 4, and 5.

6.1 CROSS-COVARIANCES

The FedGC framework intertwines data y_m^t , client state $\hat{h}_{m,a}^t$, client model parameter v_m^t , and server model parameter a_{mn}^t , creating four essential cross-covariances: Ω_m^t , Λ_m^t , Γ_{mn}^t and Ψ_{mn}^t . This section provides explicit recursions so they can be propagated together with the individual variances. From Table 1, we know that $v_m^t = \text{Vec}(\theta_m^t)$, $\mu_{\theta_m}^t = \mathbb{E}[v_m^t]$, $\Sigma_{\theta_m}^t = \text{Var}(v_m^t)$, $\mu_{y_m}^t = \mathbb{E}[y_m^t]$ and $\Omega_m^t = \text{Cov}(v_m^t, y_m^t)$. Using the above notations, Proposition 6.1 shows that v_m^t , and y_m^t are dependent with a non-zero cross-covariance Ω_m^t .

Proposition 6.1 (Client Model-Client Data Dependence). *Assume $\text{Var}(y_m^{t-1}) > 0$. Then under the federated Granger-causality updates, $\Omega_m^t := \text{Cov}(v_m^t, y_m^t) \neq 0$.*

Using Eq 4 we know that client states $\hat{h}_{m,a}^t$ are a function of the client data y_m^t . Since $\Omega_m^t \neq 0$, there must exist a dependence between the client model and client states. Proposition 6.2 analyzes the evolution of the cross-covariance between the client’s model parameter v_m^t , and the states $\hat{h}_{m,a}^t$.

Proposition 6.2 (Client Model-Client State Dependence). *Let $\Lambda_m^t = \text{Cov}(v_m^t, \hat{h}_{m,a}^t)$. Then we have the following recursion within the client, $\Lambda_m^t = \Sigma_{\theta_m}^t (I_{d_m} \otimes \mu_{y_m}^t) + \Omega_m^t (\mu_{v_m}^t \otimes I_{d_m})$,*

Due to the iterative communication between client and server, the client model dynamics are coupled with that of the server model in a feedback loop. Essentially, the client’s noisy state estimates $\hat{h}_{m,a}^t$ affect the server estimate a_{mn}^t , and the server’s uncertain a_{mn}^t in turn influences subsequent client state estimates. This effect is captured as the cross-covariance term Γ_{mn}^t given in Lemma 6.3.

Lemma 6.3 (Client State-Server Model Dependence). *The cross-covariance term $\Gamma_{mn}^t := \text{Cov}(a_{mn}^t, \hat{h}_{m,a}^t)$ follows the recursive equation: $\Gamma_{mn}^{t+1} = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t$ where, $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$, $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$, and $\Sigma_{h_m}^t = \text{Var}(\hat{h}_{m,a}^t)$.*

Because the client parameter and its augmented state are already linked through the cross-covariance in Proposition 6.2, the client state-to-server model coupling of Lemma 6.3 propagates that link one step further, yielding a direct client model-to-server model dependence captured in Lemma 6.4.

Lemma 6.4 (Client Model-Server Model Dependence). *The term $\Psi_{mn}^t := \text{Cov}(a_{mn}^t, v_m^t)$ evolves as, $\Psi_{mn}^{t+1} = D_n^t \Psi_{mn}^t H_m^{t\top} + D_n^t \Gamma_{mn}^t G_m^{t\top} - D_n^t \Sigma_{A_{mn}}^t P_m^{t\top} + 2\gamma B_{mn}^t \Lambda_m^t H_m^{t\top} + 2\gamma B_{mn}^t \Sigma_{h_m}^t G_m^{t\top} - 2\gamma B_{mn}^t \Gamma_{mn}^{t\top} P_m^{t\top}$, with the following gain matrices, $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$,*

$$D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I, G_m^t = 2\eta_1 (y_m^t \otimes (C_{mm} A_{mm})^\top), P_m^t = -2\eta_2 (y_m^t \otimes A_{mm}^\top)$$

$$\text{and } H_m^t = I_{p_m d_m} - 2\eta_1 (y_m^t y_m^{t\top}) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (y_m^t y_m^{t\top}) \otimes (A_{mm}^\top A_{mm})$$

6.2 DURING COMMUNICATION

This section characterizes the communication channel as the conduit through which every existing uncertainty—client-side data noise, client-parameter variance, and server-parameter variance—is redistributed at each iteration. Specifically we analyze the uncertainty propagation in both **(I) client-to-server**, and **(II) server-to-client** communication channel.

(I) Client to Server. At iteration t , client m sends its augmented states $\hat{h}_{m,a}^t$ to the server. While $\hat{h}_{m,a}^t$ naturally captures $\Sigma_{\theta_m}^t$ and $\Sigma_{y_m}^t$, it may also include cross-covariance $\Omega_m^t := \text{Cov}(v_m^t, y_m^t)$. Lemma 6.5 provides a closed-form expression for the uncertainty in $\hat{h}_{m,a}^t$ using $\Sigma_{\theta_m}^t$, $\Sigma_{y_m}^t$, and Ω_m^t .

Lemma 6.5 (Uncertainty in Client to Server Communication). *Let $\kappa_m = \text{tr}(\Sigma_{y_m}^t) + \|\mu_{y_m}^t\|^2$. Then the variance in the $\hat{h}_{m,a}^t$ is, $\Sigma_{h_m}^t = \kappa_m \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}$.*

(II) Server to Client. At iteration t , the server’s uncertainty is encoded in the random matrix \hat{A}_{mn}^t . Instead of sending \hat{A}_{mn}^t , the server computes and transmits the gradient: $g_{m,s}^{t+1} := \nabla_{\hat{h}_{m,a}^t} L_s^t$ where $L_s^t = \|A_{mm}[\hat{h}_{m,a}^t - \hat{h}_{m,c}^t] - \sum_{n \neq m} \hat{A}_{mn}^t \hat{h}_{n,c}^t\|^2$. This gradient inherits uncertainty from both \hat{A}_{mn}^t and $\hat{h}_{m,a}^t$, propagating the server’s model uncertainty to client m . Lemma 6.6 shows that $g_{m,s}^t$ captures the uncertainty in the server parameters, client states, and their cross-covariance.

Lemma 6.6 (Uncertainty in Server to Client Communication). *With notation as above, the uncertainty in the gradient communicated by the server is given by, $\text{Var}(g_{m,s}^{t+1}) = A_{mm}^\top U^t A_{mm}$, where $U^t = A_{mm} \Sigma_{h_m}^t A_{mm}^\top + \sum_{n \neq m} (h_{n,c}^t h_{n,c}^{t\top}) \Sigma_{A_{mn}}^t - 2 \sum_{n \neq m} A_{mm} \Gamma_{mn}^t h_{n,c}^{t\top}$.*

6.3 WITHIN SERVER

In Sections 6.1 and 6.2, we quantified how **(i)** the client-server cross-covariance Γ_{mn}^t , and **(ii)** client m ’s state variance $\Sigma_{h_m}^t$ propagate during the iterative optimization of the FedGC framework. We now analyze their contribution to the propagation of the server’s parameter uncertainty $\Sigma_{A_{mn}}^t$. Theorem 6.7 combines these components into a closed-form recursion for $\Sigma_{A_{mn}}^t$ within the server.

Theorem 6.7 (Uncertainty Propagation within the Server). *The server model parameter a_{mn}^t 's covariance evolves as, $\Sigma_{A_{mn}}^{t+1} = D_n^t \Sigma_{A_{mn}}^t D_n^{t\top} + 4\gamma^2 (\hat{h}_{n,c}^t \otimes A_{mm}) \Sigma_{h_m}^t (\hat{h}_{n,c}^t \otimes A_{mm})^\top + 2\gamma (D_n^t \Gamma_{mn}^t B_{mn}^{t\top} + B_{mn}^t \Gamma_{mn}^{t\top} D_n^{t\top})$ with $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$, and $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$*

6.4 WITHIN CLIENT

Each of the four cross-covariances discussed in Section 6.1 contribute to the propagation of uncertainty of the client model parameter θ_m (or, v_m in vectorized form). Theorem 6.8 expresses the evolution of client model's variance $\Sigma_{\theta_m}^t$ in terms of the uncertainty in its states $\Sigma_{h_m}^{t-1}$, the server model $\Sigma_{A_{mn}}^{t-1}$, and those cross-covariances Ω_m^{t-1} , Γ_{mn}^{t-1} , Ψ_{mn}^{t-1} and Λ_m^{t-1} .

Theorem 6.8 (Uncertainty Propagation within the Client). *The client-parameter covariance $\Sigma_{\theta_m}^t$ obeys the following recursion: $\Sigma_{\theta_m}^t = H_m^{t-1} \Sigma_{\theta_m}^{t-1} H_m^{t-1\top} + G_m^{t-1} \Sigma_{h_m}^{t-1} G_m^{t-1\top} + (X_m + X_m^\top) - \sum_{n \neq m} (Y_{mn} + Y_{mn}^\top) - \sum_{n \neq m} (Z_{mn} + Z_{mn}^\top) + \sum_{n \neq m} P_m^{t-1} \Sigma_{A_{mn}}^{t-1} P_m^{t-1\top}$,*

where, $X_m = H_m^{t-1} \Lambda_m^{t-1} G_m^{t-1\top}$, $Y_{mn} = H_m^{t-1} \Psi_{mn}^{t-1} P_m^{t-1\top}$, $Z_{mn} = G_m^{t-1} \Gamma_{mn}^{t-1} P_m^{t-1\top}$, and, $G_m^{t-1} := 2\eta_1 (y_m^{t-1} \otimes (C_{mm} A_{mm})^\top)$, $P_m^{t-1} := -2\eta_2 (y_m^{t-1} \otimes A_{mm}^\top)$, $H_m^{t-1} := I_{p_m d_m} - 2\eta_1 (y_m^{t-1} y_m^{t-1\top}) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (y_m^{t-1} y_m^{t-1\top}) \otimes (A_{mm}^\top A_{mm})$

7 STEADY-STATE IMPACT OF UNCERTAINTY

For tractability, we assume, **(I)** $\lim_{t \rightarrow \infty} y_m^t = \mu_{y_m}$, and **(II)** $\lim_{t \rightarrow \infty} \hat{h}_{m,c}^t = \hat{h}_{m,c} \forall m$

Proposition 7.1 (Gain Matrices Convergence). *Under the above assumptions, the gain matrices used in Section 6 converges as, $\lim_{t \rightarrow \infty} (D_n^t, H_m^t, G_m^t, P_m^t) = (D_n, H_m, G_m, P_m)$ where, $D_n = (I - 2\gamma \hat{h}_{n,c} \hat{h}_{n,c}^\top) \otimes I$, $G_m = 2\eta_1 (\mu_{y_m} \otimes (C_{mm} A_{mm})^\top)$, $P_m = -2\eta_2 (\mu_{y_m} \otimes A_{mm}^\top)$ and $H_m = I_{p_m d_m} - 2\eta_1 (\mu_{y_m} \mu_{y_m}^\top) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (\mu_{y_m} \mu_{y_m}^\top) \otimes (A_{mm}^\top A_{mm})$.*

Proposition 7.1 proves that the gains $(D_n^t, H_m^t, G_m^t, P_m^t)$ converge in the limit as $t \rightarrow \infty$.

Proposition 7.2. *If $\rho(D_n), \rho(H_m) < 1$ then we have, $\lim_{t \rightarrow \infty} (\Gamma_{mn}^t, \Psi_{mn}^t) = (\Gamma_{mn}^\infty, \Psi_{mn}^\infty)$ with $\Gamma_{mn}^\infty = (I - D_n)^{-1} 2\gamma B_{mn} \Sigma_{h_m}^\infty$, & $\Psi_{mn}^\infty = (I - H_m \otimes D_n)^{-1} \text{Vec}(D_n \Gamma_{mn}^\infty G_m^\top - D_n \Sigma_{A_{mn}}^\infty P_m^\top)$.*

With stable gains, Proposition 7.2 shows that the cross-covariance terms Γ_{mn}^t and Ψ_{mn}^t also converge, each given in closed form. Corollary 7.3 then expresses the client-state variance $\Sigma_{h_m}^\infty$ in terms of the client-parameter variance $\Sigma_{\theta_m}^\infty$ and the data moments; no other stochastic quantity survives.

Corollary 7.3. *The above assumptions lead to convergence of the uncertainty of the client states $\Sigma_{h_m}^t$ as follows: $\lim_{t \rightarrow \infty} \Sigma_{h_m}^t := \Sigma_{h_m}^\infty = \kappa_m \Sigma_{\theta_m}^\infty + \Omega_m^\infty (\mu_{y_m} \otimes I)^\top + (\mu_{y_m} \otimes I) \Omega_m^{\infty\top}$, where $\kappa_m = \text{tr}(\Sigma_{y_m}) + \|\mu_{y_m}\|^2$ and $\Omega_m^\infty = \Sigma_{\theta_m}^\infty \mu_{y_m}$.*

The key theoretical result on the impact of uncertainty quantification is mentioned next in Theorems 7.4, and 7.5, where we show that the steady-state uncertainties of server and client models are dependent only on the client data distribution (aleatoric uncertainty), and independent of the prior distribution of the parameters (epistemic uncertainty).

We know that the steady distribution of the client m 's raw data is given by $E[y_m] = \mu_{y_m}$, and $\text{Var}(y_m) = \Sigma_{y_m}$. Using $(\mu_{y_m}, \Sigma_{y_m})$ we define the following two terms that will be used in the Theorems 7.4, and 7.5: **(I)** $M_m = \mu_{y_m} \otimes I_{p_m}$, and **(II)** $\kappa_m = \text{tr}(\Sigma_{y_m}) + \|\mu_{y_m}\|^2$

Theorem 7.4 (Convergence of Server Model's Uncertainty). *Let $\rho(D_n) < 1$. Define $\mathcal{L}_n(X) := D_n X D_n^\top$ and $Q_{mn}(\Sigma) := 4\gamma^2 B_{mn} (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) B_{mn}^\top$. Then, $\Sigma_{A_{mn}}^\infty := \lim_{t \rightarrow \infty} \Sigma_{A_{mn}}^t$ exists, is unique, and is given by, $\Sigma_{A_{mn}}^\infty = \sum_{k=0}^{\infty} \mathcal{L}_n^k(Q_{mn}(\Sigma_{\theta_m}^\infty))$*

Theorem 7.5 (Convergence of Client Model's Uncertainty). *Let $\rho(H_m) < 1$. Write $\mathcal{M}_m(\Sigma) = H_m \Sigma H_m^\top$ and $R_m(\Sigma) = G_m (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) G_m^\top$. Then the steady-state $\Sigma_{\theta_m}^\infty := \lim_{t \rightarrow \infty} \Sigma_{\theta_m}^t$ is the unique solution to $\Sigma_{\theta_m}^\infty = \mathcal{M}_m(\Sigma_{\theta_m}^\infty) + R_m(\Sigma_{\theta_m}^\infty) + P_m \Sigma_{A_{mn}}^\infty P_m^\top$.*

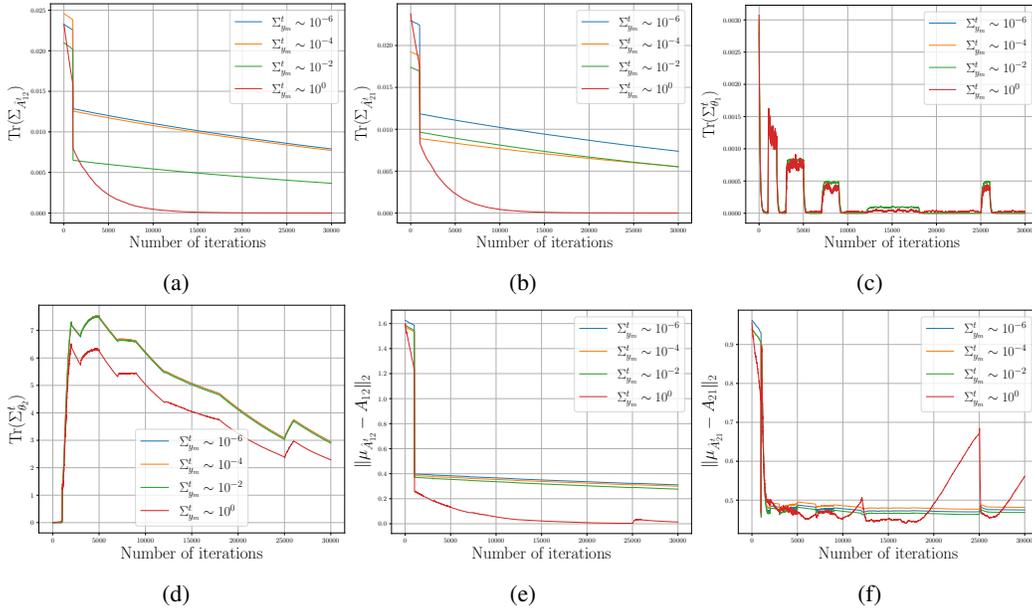


Figure 1: Uncertainty propagation during training for different levels of $\Sigma_{y_m}^t$ highlighting (a) $\text{Tr}(\Sigma_{\hat{A}_{12}}^t)$, (b) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (c) $\text{Tr}(\Sigma_{\theta_1}^t)$, (d) $\text{Tr}(\Sigma_{\theta_2}^t)$, (e) $\|\mu_{\hat{A}_{12}}^t - A_{12}\|_2$, (f) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$ vs iterations.

Theorem 7.4 shows that the steady state uncertainty of the server parameter represented by $\Sigma_{A_{mn}}^\infty$ depends only on the distribution of client data $(\mu_{y_m}, \Sigma_{y_m})$, and the steady state client model’s variance $\Sigma_{\theta_m}^\infty$. It is **independent of the prior variance** $\Sigma_{\theta_m}^0$. Similarly, Theorem 7.5 establishes that the steady-state variance of the client model parameter $\Sigma_{\theta_m}^\infty$ is uniquely determined by the client data distribution $(\mu_{y_m}, \Sigma_{y_m})$, and the converged server uncertainty $\Sigma_{A_{mn}}^\infty$. Crucially, this result confirms that the client’s epistemic uncertainty is governed entirely by aleatoric quantities and training dynamics, and is **independent of the initial variance** $\Sigma_{\theta_m}^0$.

8 EXPERIMENTS

8.1 SYNTHETIC DATASET

Experimental Settings. We simulate a multi-client linear time-invariant (LTI) state space model described in 3. To enable interpretability, the experiments (except scalability) rely on a two-client setup with $p_m = 2$ and $d_m = 8$ for $m = \{1, 2\}$ such that the off-diagonal blocks of the state matrix (A) are $A_{12} = 0$ and $A_{21} \neq 0$. Both client and server models are regularized to ensure feasible solutions. Further experimental details are provided in Section A of the Appendix.

Aleatoric. To analyze the effects of aleatoric noise, we change the data variance Σ_{y_m} and observe the uncertainty of the server and client models at each iteration. The trace of covariance of the server parameter ($\hat{A}_{12}^t, \hat{A}_{21}^t$) and the corresponding parameters (θ_1^t, θ_2^t) at clients 1 and 2, respectively are plotted for different Σ_{y_m} in Figures 1(a)-(f). The evolution plots for the covariance of both server and client models validate the claims (in Section 6, and Theorems 7.4 & 7.5).

The jumps in the trace covariance plots of (θ_1^t, θ_2^t) refer to the points where the mean-shifts of data (i.e., μ_{y_m}) occur. A preliminary inspection of Figures 1(a)-(e) shows that higher Σ_{y_m} accelerates variance decay and hence faster learning. However, we can see from the norm error plots in Figure 1(f) that very high Σ_{y_m} causes larger estimation error in learning, and this effect is more profound during mean shifts. These errors stem from violating Assumption (A4) (Section 5) in mean-shifts.

Epistemic. The effect of epistemic uncertainty is demonstrated by sampling \hat{A}_{mn}^0 and θ_m^0 from normal distributions while progressively increasing their variance. This allows us to observe how uncertainty propagates through the server and client models at each iteration. Figure 2 and Figure 3

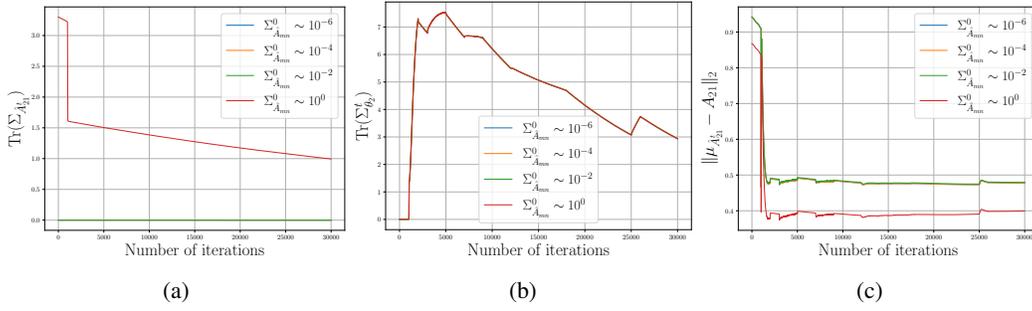


Figure 2: Uncertainty propagation during training for different levels of $\Sigma_{\hat{A}_{mn}}^0$ highlighting (a) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (b) $\text{Tr}(\Sigma_{\theta_2}^t)$, (c) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$ vs iterations.

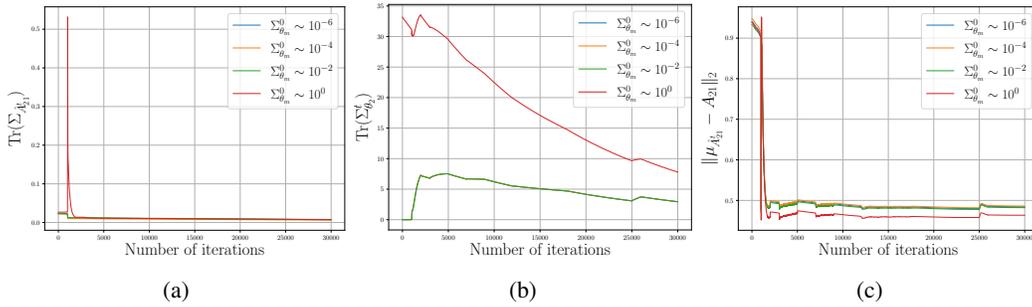


Figure 3: Uncertainty propagation during training for different levels of $\Sigma_{\theta_m}^0$ highlighting (a) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (b) $\text{Tr}(\Sigma_{\theta_2}^t)$, (c) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$ vs iterations.

shows that the evolution of the covariances for server and client models is agnostic to the variance of the prior normal distribution, further validating the claim of Theorems 7.4 and 7.5.

Please refer to Appendix B for additional results highlighting the evolution of the (a) cross-covariance terms, (b) communicated terms (mentioned in Section 6), and (c) Scalability studies.

8.2 REAL-WORLD DATASET

We perform real-world experiments on two industrial cybersecurity datasets: (1) HAI Shin et al. (2021) and (2) SWaT Mathur & Tippenhauer (2016), each containing multiple interacting processes. We fit a linear state space model separately to each process. Specifically, we use the subspace identification method Overschee & Moor (1994) to identify the process-specific state-transition and observation matrices, which are further used for analyzing theoretical results of Section 6 & 7. Preprocessing details are provided in Appendix C.

Results. Cross-process GC in both datasets is learned using the FedGC framework. Owing to space constraints, we provide the uncertainty evolution plots and the convergence results in Appendix C.

9 LIMITATIONS

The analysis is based on the FedGC framework, which can be restrictive due to its linear models. We provide extensions to non-linear models in Appendix E. The uncertainty propagation relies on data stationarity, and we discuss possible relaxations in Appendix E.2, though a full treatment remains open. Theoretical guarantees scale unfavorably with dimensionality, which limits applicability to very large systems; we complement this by analyzing both communication and computation complexity in Appendix D.5. Finally, client-server cross-covariance terms may expose sensitive data; we present preliminary differential privacy modeling in Appendix D.4.

REFERENCES

- 486
487
488 Marc Jayson Baucas and Petros Spachos. Federated kalman filter for secure iot-based device moni-
489 toring services. *IEEE Networking Letters*, 5(2):91–94, 2023.
- 490 Matthew Chan, Maria Molina, and Chris Metzler. Estimating epistemic and aleatoric uncertainty
491 with a single model. *Advances in Neural Information Processing Systems*, 37:109845–109870,
492 2024.
- 493 Hong-You Chen and Wei-Lun Chao. Fedbe: Making bayesian model ensemble applicable to feder-
494 ated learning. *arXiv preprint arXiv:2009.01974*, 2020.
- 496 Nikita Durasov, Doruk Oner, Jonathan Donier, Hieu Le, and Pascal Fua. Enabling uncertainty
497 estimation in iterative neural networks. *arXiv preprint arXiv:2403.16732*, 2024.
- 498 Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt,
499 Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey
500 of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589,
501 2023.
- 503 Han Guo, Philip Greengard, Hongyi Wang, Andrew Gelman, Yoon Kim, and Eric Xing. Federated
504 learning as variational inference: A scalable expectation propagation approach. In *The Eleventh
505 International Conference on Learning Representations*.
- 506 Xianjie Guo, Liping Yi, Xiaohu Wu, Kui Yu, and Gang Wang. Enhancing causal discovery in fed-
507 erated settings with limited local samples. In *International Workshop on Trustworthy Federated
508 Learning*, pp. 164–179. Springer, 2024a.
- 509 Xianjie Guo, Kui Yu, Lin Liu, and Jiuyong Li. Fedcsl: a scalable and accurate approach to feder-
510 ated causal structure learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
511 volume 38, pp. 12235–12243, 2024b.
- 513 Paul Hofman, Yusuf Sale, and Eyke Hüllermeier. Quantifying aleatoric and epistemic uncertainty: A
514 credal approach. In *ICML 2024 Workshop on Structured Probabilistic Inference {\&} Generative
515 Modeling*, 2024.
- 516 Eyke Hüllermeier and Willem Waegeman. Aleatoric and epistemic uncertainty in machine learning:
517 An introduction to concepts and methods. *Machine learning*, 110(3):457–506, 2021.
- 518 Denis Huseljic, Bernhard Sick, Marek Herde, and Daniel Kottke. Separation of aleatoric and epis-
519 temic uncertainty in deterministic deep neural networks. In *2020 25th International Conference
520 on Pattern Recognition (ICPR)*, pp. 9172–9179. IEEE, 2021.
- 522 George N Korres. A distributed multiarea state estimation. *IEEE Transactions on Power Systems*,
523 26(1):73–84, 2010.
- 524 Nikita Kotelevskii, Maxime Vono, Alain Durmus, and Eric Moulines. Fedpop: A bayesian approach
525 for personalised federated learning. *Advances in Neural Information Processing Systems*, 35:
526 8687–8701, 2022.
- 528 Shenghui Li, Edith Ngai, and Thiemo Voigt. Byzantine-robust aggregation in federated learning
529 empowered industrial iot. *IEEE Transactions on Industrial Informatics*, 19(2):1165–1175, 2021.
- 530 Aditya P. Mathur and Nils Ole Tippenhauer. Swat: a water treatment testbed for research and
531 training on ics security. In *2016 International Workshop on Cyber-physical Systems for Smart
532 Water Networks (CySWater)*, pp. 31–36, 2016.
- 534 Nis Meinert, Jakob Gawlikowski, and Alexander Lavin. The unreasonable effectiveness of deep ev-
535 idential regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37,
536 pp. 9134–9142, 2023.
- 537 Ayush Mohanty, Nazal Mohamed, Paritosh Ramanan, and Nagi Gebraeel. Federated granger causal-
538 ity learning for interdependent clients with state space representation. *International conference
539 on learning representations*, 2025.

- 540 A Muir and J Lopatto. Final report on the august 14, 2003 blackout in the united states and canada:
541 causes and recommendations. 2004.
542
- 543 Peter Van Overschee and Bart De Moor. N4sid: Subspace algorithms for the identification of
544 combined deterministic-stochastic systems. *Autom.*, 30:75–93, 1994. URL <https://api.semanticscholar.org/CorpusID:28586805>.
545
- 546 Junha Park, Jiseon Moon, Taekyoon Kim, Peng Wu, Tales Imbiriba, Pau Closas, and Sunwoo Kim.
547 Federated learning for indoor localization via model reliability with dropout. *IEEE Communica-*
548 *tions Letters*, 26(7):1553–1557, 2022.
549
- 550 Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning.
551 *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022.
- 552 Anggoro Primadianto and Chan-Nan Lu. A review on distribution system state estimation. *IEEE*
553 *Transactions on Power Systems*, 32(5):3875–3883, 2016.
554
- 555 Hyeok-Ki Shin, Woomyo Lee, Jeong-Han Yun, and Byung-Gi Min. Two ics security datasets
556 and anomaly detection contest on the hil-based augmented ics testbed. In *Cyber Security Ex-*
557 *perimentation and Test Workshop*, CSET ’21, pp. 36–40, New York, NY, USA, 2021. Associa-
558 tion for Computing Machinery. ISBN 9781450390651. doi: 10.1145/3474718.3474719. URL
559 <https://doi.org/10.1145/3474718.3474719>.
- 560 Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of*
561 *Statistics and Its Application*, 9(1):289–319, 2022.
- 562 Alex Tank, Ian Covert, Nicholas Foti, Ali Shojaie, and Emily B Fox. Neural granger causality. *IEEE*
563 *Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4267–4279, 2021.
564
- 565 Florian Van Daalen, Lianne Ippel, Andre Dekker, and Inigo Bermejo. Vertibayes: learning bayesian
566 network parameters from vertically partitioned data with missing values. *Complex & Intelligent*
567 *Systems*, 10(4):5317–5329, 2024.
- 568 Elahe Vedadi, Joshua V Dillon, Philip Andrew Mansfield, Karan Singhal, Arash Afkanpour, and
569 Warren Richard Morningstar. Federated variational inference: Towards improved personalization
570 and generalization. In *Proceedings of the AAAI Symposium Series*, volume 3, pp. 323–327, 2024.
571
- 572 Hanjing Wang, Dhiraj Joshi, Shiqiang Wang, and Qiang Ji. Gradient-based uncertainty attribution
573 for explainable bayesian deep learning. In *Proceedings of the IEEE/CVF Conference on Computer*
574 *Vision and Pattern Recognition*, pp. 12044–12053, 2023.
- 575 Axel Wismüller, Adora M Dsouza, M Ali Vosoughi, and Anas Abidin. Large-scale nonlinear granger
576 causality for inferring directed dependence from short multivariate time-series data. *Scientific*
577 *reports*, 11(1):7817, 2021.
- 578 Zirui Xing and Yuanqing Xia. Distributed federated kalman filter fusion over multi-sensor unreliable
579 networked systems. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 63(10):1714–
580 1725, 2016.
- 581 Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept
582 and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19,
583 2019.
- 584
585 Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and
586 Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *Internat-*
587 *ional conference on machine learning*, pp. 7252–7261. PMLR, 2019.
588
- 589 Xu Zhang, Yinchuan Li, Wenpeng Li, Kaiyang Guo, and Yunfeng Shao. Personalized federated
590 learning via variational bayesian inference. In *International Conference on Machine Learning*,
591 pp. 26293–26310. PMLR, 2022.
592
593

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

Appendices

CONTENTS

A	The Use of Large Language Models (LLMs)	14
B	Additional Results on on Synthetic Dataset	14
	B.1 Experimental Details	14
	B.1.1 Cross-Covariance and Communicated Terms	14
	B.2 Scalability Studies	14
	B.3 Comparison to Baselines	15
	B.4 Effect of DP Noise on Uncertainty and Causal Accuracy	15
C	Results on Real-World Dataset	16
	C.1 Experimental Setting	16
	C.2 Preprocessing	17
	C.3 Results	17
D	Additional Theoretical Results	17
	D.1 Lemma D.1	17
	D.2 Proposition D.2	18
	D.3 Proposition D.3	18
	D.4 Privacy Analysis	19
	D.5 Complexity Analysis	20
E	Extensions to the Framework	20
	E.1 Non-Linear Models	20
	E.2 Relaxing Stationary Assumptions	22
F	Proofs	22
	F.1 Proposition 6.1	22
	F.2 Proposition 6.2	23
	F.3 Lemma 6.3	24
	F.4 Lemma 6.4	24
	F.5 Lemma 6.5	25
	F.6 Lemma 6.6	26
	F.7 Theorem 6.7	26
	F.8 Theorem 6.8	27
	F.9 Proposition 7.1	28
	F.10 Proposition 7.2	29

648	F.11 Corollary 7.3	30
649	F.12 Theorem 7.4	30
650	F.13 Theorem 7.5	31
651		
652		
653		
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		

Note. All the plots corresponding to Section B, and C are present at the end of this document.

A THE USE OF LARGE LANGUAGE MODELS (LLMs)

In preparing this work, large language models (LLMs), specifically ChatGPT, were used only as a writing aid to polish the clarity, flow, and readability of the manuscript. LLMs were not involved in generating or shaping the underlying theoretical ideas, arguments, or results. All core contributions, analyses, and conclusions presented here are entirely original and solely attributable to the authors.

B ADDITIONAL RESULTS ON ON SYNTHETIC DATASET

B.1 EXPERIMENTAL DETAILS

Experiments on synthetic dataset was conducted by simulating a two-client LTI system. Each client has hidden latent state dimension $p_m = 2$ and measurement dimension $d_m = 8$. A one way dependency in which client 1 granger causes client 2 (and not vice-versa) is considered. Therefore $A_{12} = 0$ and $A_{21} \neq 0$. The loss functions in (5) and (8) are regularized and the learning rates γ, η_1 and η_2 are adjusted to ensure convergence but not optimally tuned.

All the results reported are in terms of the trace of the covariance matrices. Trace of the covariance matrix was chosen as a scalar quantity to quantify the uncertainty from the covariance matrices for a multitude of reasons such as: 1) it is the sum of variances across all directions, 2) it is rotation invariant, 3) computationally cheap etc. Unless otherwise specified, the quantities plotted are relevant to explaining the uncertainty propagation in learning \hat{A}_{21} .

As discussed in Section 8, the effect of aleatoric noise in all the experiments is studied by changing the data variance $\Sigma_{y_m}^t$ and the effect of epistemic noise is studied by changing the variance of the initial values $\Sigma_{\hat{A}_{mn}}^0$ and $\Sigma_{\theta_m}^t$.

B.1.1 CROSS-COVARIANCE AND COMMUNICATED TERMS

Cross-Covariances. The evolution of cross-covariance terms $\Lambda_{m,n}^t, \Gamma_{m,n}^t$ and $\Psi_{m,n}^t$ discussed in Section 6 for different regimes of aleatoric and epistemic noise are given in Figures 5, 6 and 7. We can observe that the cross-covariance terms converge even for very high noise regimes.

During Communication. The augmented client state $\hat{h}_{m,a}^t$ sent from the client carries the uncertainty from the client to the server and the gradient $\nabla_{\hat{h}_{m,a}^t} L_s^t$ sent back to the client carries the uncertainty from the server to the client. The evolution of covariances of these communicated terms for different regimes of aleatoric and epistemic noises are plotted in Figures 8, 9 and 10.

B.2 SCALABILITY STUDIES

In the scalability experiment, we consider the two-client system with similar causal relationship as in previous experiments. The hidden state dimension $p_m = 2$ is kept constant while the measurement dimension d_m is increased for both clients. Trace of covariance of \hat{A}_{21} at convergence for regimes of $\Sigma_{\hat{A}_{mn}}^0 \sim \{10^{-6}, 10^{-4}, 10^{-2}, 10^0\}$ is summarized in Table 2. We could observe that the performance of the framework remains more or less similar with increased dimension of the measurements d_m .

Table 2: Trace(Cov(\hat{A}_{21})) vs measurement (i.e., raw data) dim. d_m

Order of Variance	Measurement (Raw Data) Dimension (d_m)			
	$d_m = 16$	$d_m = 32$	$d_m = 64$	$d_m = 128$
$\sim 10^{-6}$	$\approx 10^{-9}$	$\approx 10^{-5}$	$\approx 10^{-7}$	$\approx 10^{-7}$
$\sim 10^{-4}$	$\approx 10^{-8}$	$\approx 10^{-5}$	$\approx 10^{-7}$	$\approx 10^{-7}$
$\sim 10^{-2}$	1.0×10^{-4}	4.0×10^{-4}	1.0×10^{-4}	$\approx 10^{-5}$
$\sim 10^0$	1.6722	1.9733	2.4601	1.5882

In the second experiment, we keep $p_m = 2$ and $d_m = 8$ and vary the number of clients M . We report the trace of covariance of \hat{A}_{21} at convergence for different regimes of $\Sigma_{\hat{A}_{mn}}^0$ in Table 3.

Table 3: Trace(Cov(\hat{A}_{21})) vs. no. of clients M

Order of Variance	Number of Clients (M)			
	$M = 2$	$M = 4$	$M = 8$	$M = 16$
$\sim 10^{-6}$	$\approx 10^{-5}$	$\approx 10^{-5}$	$\approx 10^{-5}$	$\approx 10^{-6}$
$\sim 10^{-4}$	$\approx 10^{-5}$	$\approx 10^{-5}$	$\approx 10^{-5}$	$\approx 10^{-6}$
$\sim 10^{-2}$	0.0001	0.0002	0.0002	0
$\sim 10^0$	1.7233	1.6456	2.5835	4.7

B.3 COMPARISON TO BASELINES

To contextualize our uncertainty-propagation framework, we compare our method against several natural baselines for vertically federated Granger-causal learning. Table 4 reports two metrics: **(i)** the Frobenius error $\|\hat{A}_{21} - A_{21,\text{true}}\|_F$ measuring estimation accuracy, and **(ii)** the trace $\text{tr}(\Sigma_{A_{21}})$ of the propagated covariance, quantifying parameter uncertainty.

- **Centralized.** All client data are pooled and A_{21} is estimated in a non-federated, oracle fashion. This constitutes the ideal upper bound on estimation accuracy because no privacy or communication constraints are present.
- **Our Method** Our method computes empirical uncertainty by running FedGC from *multiple independently sampled initializations* (multi-start ensemble). Each run follows the standard FedGC alternating client-server updates. The empirical mean and covariance of the resulting estimates are reported.
- **Independent Clients.** Each client estimates its own local block without modeling cross-client dependencies (equivalent to setting $\eta_2 = 0$ in Eq. (5)). This baseline tests the effect of removing inter-client coupling altogether.

Table 4 summarizes the comparison. As expected, Centralized estimation achieves the lowest error and nearly zero uncertainty since it has full access to all data. Our method achieves the best performance among federated approaches: it produces the lowest Frobenius error and the smallest covariance trace, indicating both accurate and stable estimation of A_{21} . In contrast, Independent Clients baseline underperform: removing iterative coupling leads to a higher estimation error and significantly larger propagated uncertainty.

Table 4: Comparison across baselines.

Method	$\ \hat{A}_{21} - A_{21,\text{true}}\ _F$	$\text{tr}(\Sigma_{A_{21}})$
Centralized	4.804058×10^{-1}	3.960954×10^{-14}
Our	4.96649×10^{-1}	7.96647×10^{-9}
Independent Clients	5.02350×10^{-1}	1.87971×10^{-8}

B.4 EFFECT OF DP NOISE ON UNCERTAINTY AND CAUSAL ACCURACY

In this subsection, we study how differentially private (DP) Gaussian noise injected into the federated messages affects **(i)** steady-state parameter uncertainty and **(ii)** accuracy of causal link detection. We consider two communication directions: *client*→*server*, where noise is added to the states sent by the clients to the server, and *server*→*client*, where noise is added to the gradients sent back to the clients. In both cases, we use the Gaussian mechanism and sweep the noise standard deviation over a log-scale grid $\sigma \in \{10^{-6}, 10^{-5}, \dots, 10^{-1}\}$.

Figure 4 summarizes the results on the synthetic dataset. The top row shows the steady-state uncertainty of the cross-client Granger block A_{21} , measured via $\text{tr}(\Sigma_{A_{21}})$, as a function of the in-

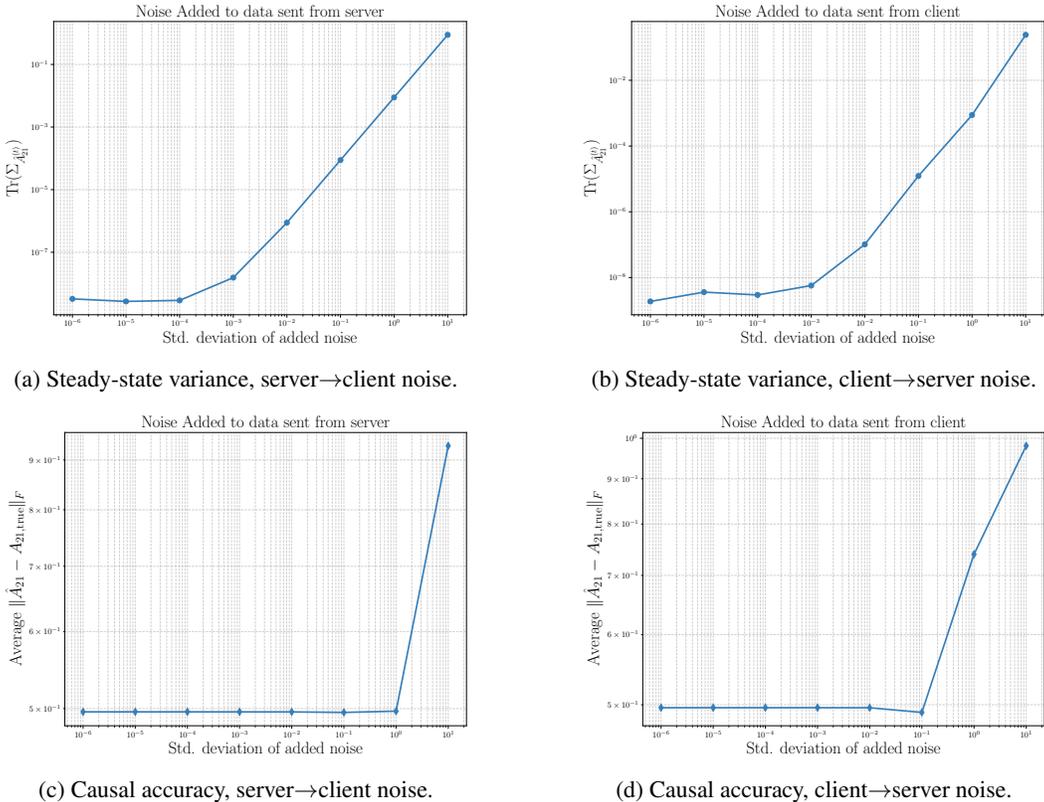


Figure 4: Effect of DP Gaussian noise on steady-state uncertainty and causal link detection in the synthetic dataset. Top row: trace of the propagated covariance $\text{tr}(\Sigma_{A_{21}})$ as a function of the noise scale σ . Bottom row: Frobenius error $\|\hat{A}_{21} - A_{21,\text{true}}\|_F$ (normalized) versus σ .

jected noise scale σ . The bottom row reports the corresponding causal estimation accuracy, measured by the Frobenius norm $\|\hat{A}_{21} - A_{21,\text{true}}\|_F$ and normalized for visualization. Panels labelled “server→client” and “client→server” correspond to the direction in which DP noise is added. Across all configurations, the uncertainty recursions remain numerically stable for the entire range of σ , and the steady-state variance increases smoothly as the DP noise level grows. For small to moderate noise ($\sigma \lesssim 10^{-3}$), the inflation in $\text{tr}(\Sigma_{A_{21}})$ is modest and the causal estimation error remains close to the non-DP baseline in both communication directions. Only for the largest noise levels ($\sigma \approx 10^{-2}$ – 10^{-1}) do we observe a pronounced increase in variance accompanied by a degradation in causal link detection, as expected from strong DP perturbations. These experiments confirm that our uncertainty propagation framework is robust to reasonably strong DP noise and that the empirical privacy–utility trade-off behaves consistently with the Gaussian mechanism in both communication channels.

C RESULTS ON REAL-WORLD DATASET

C.1 EXPERIMENTAL SETTING

We utilize two industrial cybersecurity datasets – (1) HAI and (2) SWaT, to conduct real-world experiments.

HAI. The HAI dataset captures SCADA (Supervisory Control and Data Acquisition) time-series data from a realistic industrial control system testbed augmented with a Hardware-in-the-Loop (HIL) simulator that emulates steam-turbine power generation and pumped-storage hydropower generation. HAI testbed comprises four processes: 1. Boiler Process (P1), 2. Turbine Process (P2), 3. Water Treatment (P3), 4. HIL Simulator (P4). Processes P1–P3 each represent a physical subsys-

tem with multiple sensor measurements and are treated as individual clients in our study. On the other hand, P4 is a HIL simulation engine exposing control signals rather than sensor outputs and is therefore excluded from our analysis.

SWaT. SWaT is the data from an industry-compliant six-stage water treatment testbed commissioned in March 2015 by iTrust. Different stages in water treatment includes raw water storage, pH balancing, ultrafiltration, dechlorination, reverse osmosis, and membrane backwash. Each stage is treated as a client in our setting. We do not include the last stage (membrane backwash) in our study because all the measurements in this stage are constant except one that takes few more values and inclusion of this stage was causing convergence issues.

C.2 PREPROCESSING

We utilize different preprocessing methodologies for the two real-world datasets.

HAI. We first normalized the measurements to make sure they are evenly scaled. A subspace identification algorithm was used to fit an LTI system to these normalized measurements. The subspace identification outputs the state-transition matrix A (which also signifies the Granger causality), and the observation matrix C . The size of low-dimensional states p for the entire system was obtained by looking at the decay of singular values of Hankel matrix involved in the subspace identification step.

The observation matrix C obtained from subspace identification is not block diagonal. An L_2 norm based thresholding was used to assign each state variable to a process and build a block diagonal C matrix. This is followed by re-estimating A matrix using least squares with the block diagonal C matrix. This re-estimated A matrix is considered the ground truth Granger causality for our study.

SWaT. For each client, the measurement variables with pearson correlation greater 0.3 with other clients are selected. This is followed by the same steps as in the case of HAI data.

C.3 RESULTS

We plot the evolution of covariance of all the off-diagonal blocks \hat{A}_{mn}^t , covariance of θ_m^t for all clients and average L_2 norm error of each \hat{A}_{mn}^t for both datasets in Figures 17-28.

HAI. The server-side uncertainty on all off-diagonal \hat{A}_{mn} terms drops rapidly within the first few hundred iterations and stabilizes at very low traces (Figure 17), indicating fast convergence even under high data (measurement) noise. Local parameter variances Σ_{θ_m} likewise shrink to nearly zero (Figure 6), with only minor differences between the three processes. The average estimation error $\|\mu_{A_{mn}} - A_{mn}\|$ decreases monotonically on the log-iteration scale (Figure 7), confirming that reduced uncertainty directly translates to higher causal-link accuracy.

SWaT. In contrast to HAI dataset, SWaT’s causal-block covariances decline more gradually over $\sim 1,000$ iterations and do not show signs of complete convergence (Figure 25). This suggests that the six-stage water-treatment testbed requires more iterations or additional data to fully resolve the underlying causal links. Local model variances Σ_{θ_m} exhibit a wider spread (at the end of 1000 iterations), with some retaining significantly more uncertainty than others (Figure 27). Correspondingly, the ℓ_2 estimation errors plateau earlier and remain higher (Figure 28), underscoring the slower, higher-variance convergence behavior in SWaT, possibly due to higher noise in the collected data.

D ADDITIONAL THEORETICAL RESULTS

D.1 LEMMA D.1

Lemma D.1. *At any time t and, for each client m , define*

$$g_m = \nabla_{\theta_m} L_s = X_m (y_m^{t-1})^\top, \quad X_m := A_{mm}^\top \left(A_{mm} \Delta h_m - \sum_{r \neq m} \hat{A}_{mr} (\hat{h}_r^{t-1})_c \right).$$

Under Assumptions (A1)-(A4), we have for any $m \neq n$,

$$\text{Cov}(\text{vec}(g_m), \text{vec}(g_n)) = 0.$$

918 *Proof.* By Assumption **(A2)**, the block-row server parameters for clients m and n are mutually
 919 independent, hence $X_m \perp X_n$. By Assumptions **(A1)**, **(A4)**, and **(A5)**, the local data satisfy

$$920 \mu_{y_m} = \mathbb{E}[y_m^{t-1}], \quad \Sigma_{y_m} = \text{Cov}(y_m^{t-1}, y_m^{t-1}),$$

921 with $\text{Cov}(y_m^{t-1}, y_n^{t-1}) = 0$ for $m \neq n$. Since y_m^{t-1} is also independent of $\{A_{pq}^t\}$ by Assumption
 922 **(A3)**, we may write

$$923 \text{vec}(g_m) = (I \otimes X_m) y_m^{t-1}.$$

924 Thus, for $m \neq n$,

$$925 \text{Cov}(\text{vec}(g_m), \text{vec}(g_n)) = \mathbb{E}[(I \otimes X_m) y_m^{t-1} (y_n^{t-1})^\top (I \otimes X_n)^\top] - \mathbb{E}[\text{vec}(g_m)] \mathbb{E}[\text{vec}(g_n)]^\top.$$

926 Independence allows factorization, and Assumption **(A4)**–**(A5)** implies $\mathbb{E}[y_m^{t-1} (y_n^{t-1})^\top] = \mu_{y_m} \mu_{y_n}^\top$,
 927 which cancels with the product of means. Hence

$$928 \text{Cov}(\text{vec}(g_m), \text{vec}(g_n)) = 0, \quad m \neq n.$$

929 \square

930 D.2 PROPOSITION D.2

931 **Proposition D.2.** *Consider the client update rule*

$$932 \theta_m^{t+1} = \theta_m^t - \eta_1 g_{m,a}^t - \eta_2 g_{m,s}^t,$$

933 where $g_{m,a}^t = \nabla_{\theta_m} L_m$ (local) and $g_{m,s}^t = \nabla_{\theta_m} L_s$ (global). Suppose Assumptions **(A1)**–**(A5)** hold.
 934 Then, for any two distinct clients $m \neq n$,

$$935 \lim_{t \rightarrow \infty} \text{Cov}(\theta_m^t, \theta_n^t) = 0.$$

936 *Proof.* Expanding the covariance for $m \neq n$ gives

$$937 \text{Cov}(\theta_m^{t+1}, \theta_n^{t+1}) = \text{Cov}(\theta_m^t, \theta_n^t) + \eta_2^2 \text{Cov}(g_{m,s}^t, g_{n,s}^t) + \text{vanishing local terms}$$

938 where the omitted terms involve cross-products between local and global gradients.

939 By Lemma D.1, $\text{Cov}(g_{m,s}^t, g_{n,s}^t) = 0$ for $m \neq n$. By **(A2)** and **(A3)**, $\theta_m^t \perp \theta_n^t$ at initialization. By
 940 **(A4)**–**(A5)**, client data are uncorrelated across m, n , and all noise is additive. Hence all cross-terms
 941 vanish, and we obtain the recursion

$$942 \text{Cov}(\theta_m^{t+1}, \theta_n^{t+1}) = \text{Cov}(\theta_m^t, \theta_n^t).$$

943 Since the initial covariance is zero, induction gives $\text{Cov}(\theta_m^t, \theta_n^t) = 0$ for all t , and thus

$$944 \lim_{t \rightarrow \infty} \text{Cov}(\theta_m^t, \theta_n^t) = 0.$$

945 \square

946 D.3 PROPOSITION D.3

947 As a direct consequence of Assumption **(A5)** we have the following proposition:

948 **Proposition D.3 (Data Uncertainty with Zero Process Noise).** *Consider the discrete LTI system*

$$949 x^{t+1} = A x^t + w^t, \quad y^t = C x^t + v^t, \quad w^t \sim \mathcal{N}(0, Q), \text{ and } v^t \sim \mathcal{N}(0, R)$$

950 and assume the state matrix is asymptotically stable, $\rho(A) < 1$. In the absence of process noise i.e.,
 951 $Q = 0$, then the data variance converges to its noise covariance. $\lim_{t \rightarrow \infty} \text{Var}(y^t) = R$.

952 *Proof.* With $Q = 0$, the state update becomes deterministic:

$$953 x^{t+1} = A x^t.$$

954 Thus, assuming some initial condition x^0 , the state at time t is given by,

$$955 x^t = A^t x^0.$$

Therefore, the data y^t is given by,

$$y^t = Cx^t + v^t = CA^t x^0 + v^t.$$

To compute the variance of y^t , we consider the following equation,

$$\text{Var}(y^t) = \text{Var}(CA^t x^0 + v^t).$$

Since $v^t \sim \mathcal{N}(0, R)$ is independent of x^0 , and $CA^t x^0$ is deterministic, we write $\text{Var}(y^t)$ as,

$$\text{Var}(y^t) = \text{Var}(CA^t x^0) + \text{Var}(v^t).$$

Now, since $CA^t x^0$ is deterministic, its variance is given by,

$$\text{Var}(CA^t x^0) = CA^t \text{Var}(x^0)(A^t)^\top C^\top.$$

Given that $\rho(A) < 1$, we have $A^t \rightarrow 0$ as $t \rightarrow \infty$ we have,

$$CA^t \text{Var}(x^0)(A^t)^\top C^\top \rightarrow 0.$$

Using the above result, and the fact that $\text{Var}(v^t) = R$ we can write $\text{Var}(y^t)$ as,

$$\lim_{t \rightarrow \infty} \text{Var}(y^t) = \lim_{t \rightarrow \infty} [CA^t \text{Var}(x^0)(A^t)^\top C^\top + R] = R.$$

□

D.4 PRIVACY ANALYSIS

We formally prove that the cross-covariance terms between client quantities and server quantities can be made differentially private under the Gaussian mechanism.

FedGC Latent Privacy. In the original FedGC setup, each client m communicates a compressed latent state $\hat{h}_{c,m}^t$ and an augmented latent state $\hat{h}_{a,m}^t$, both of which are computed from private data y_m^t via local encoders (KF). As shown in Appendix F of FedGC, if the mapping $y_m^t \mapsto \hat{h}_{c,m}^t$ has bounded ℓ_2 -sensitivity Δ , then adding Gaussian noise:

$$\tilde{h}_{c,m}^t = \hat{h}_{c,m}^t + \mathcal{N}(0, \sigma^2 I), \quad \sigma \geq \frac{\Delta \cdot \sqrt{2 \log(1.25/\delta)}}{\varepsilon}$$

ensures (ε, δ) -differential privacy for each client's latent state. The same construction holds for $\hat{h}_{a,m}^t$ and is preserved in our framework.

DP for Cross-Covariances. Unlike FedGC, our method introduces server-side use of *cross-covariance* matrices between client representations, which may leak client-private correlations. The key objects are:

$$\Gamma_{mn}^t := \text{Cov}(\hat{h}_{c,m}^t, \hat{h}_{a,n}^t) = \mathbb{E}[\hat{h}_{c,m}^t (\hat{h}_{a,n}^t)^\top] - \bar{h}_{c,m}^t (\bar{h}_{a,n}^t)^\top$$

$$\Psi_{mn}^t := \text{Cov}(v_m^t, \hat{h}_{c,n}^t) = \mathbb{E}[v_m^t (\hat{h}_{c,n}^t)^\top] - \bar{v}_m^t (\bar{h}_{c,n}^t)^\top$$

where $v_m^t = \text{vec}(\theta_m^t)$ is the flattened client model parameter. We now show how to make these matrices differentially private via Gaussian noise.

Each client m transmits perturbed values of $\hat{h}_{c,m}^t, \hat{h}_{a,m}^t$

$$\tilde{h}_{c,m}^t = \hat{h}_{c,m}^t + \xi_c^t, \quad \tilde{h}_{a,n}^t = \hat{h}_{a,n}^t + \xi_a^t$$

with $\xi_c^t, \xi_a^t \sim \mathcal{N}(0, \sigma_h^2 I)$. Furthermore, since cross-covariance Ψ_{mn} requires v_m , we perturb it as:

$$\tilde{v}_m^t = v_m^t + \xi_v^t$$

with $\xi_v^t \sim \mathcal{N}(0, \sigma_v^2 I)$. The server then computes the cross-covariances:

$$\tilde{\Gamma}_{mn}^t := \text{Cov}(\tilde{h}_{c,m}^t, \tilde{h}_{a,n}^t), \quad \tilde{\Psi}_{mn}^t := \text{Cov}(\tilde{v}_m^t, \tilde{h}_{c,n}^t)$$

These estimators contain the desired cross-covariance along with stochastic masking from the perturbations.

Using classical state-space theory & Lipschitz assumptions, we have norm bounds: $\|\hat{h}_{c,m}^t\|_2 \leq B_c$, $\|\hat{h}_{a,n}^t\|_2 \leq B_a$, $\|v_m^t\|_2 \leq B_v$. Then, the ℓ_2 -sensitivities of Γ_{mn} , and Ψ_{mn} are:

$$\Delta_2(\Gamma_{mn}^t) \leq 2B_c B_a, \quad \Delta_2(\Psi_{mn}^t) \leq 2B_v B_c$$

To achieve (ϵ, δ) -DP via the Gaussian mechanism, it suffices to add i.i.d. noise to each element of the matrices:

$$\sigma_\Gamma \geq \frac{2B_c B_a \sqrt{2 \log(1.25/\delta)}}{\epsilon}, \quad \sigma_\Psi \geq \frac{2B_v B_c \sqrt{2 \log(1.25/\delta)}}{\epsilon}$$

This guarantees that the server-observed $\tilde{\Gamma}_{mn}^t$ and $\tilde{\Psi}_{mn}^t$ are differentially private with respect to any one client’s data at time t .

Implications on Uncertainty Propagation. The injection of DP noise into Γ_{mn}^t and Ψ_{mn}^t affects the downstream uncertainty estimates in both client- and server-side propagation. Specifically, Γ_{mn}^t appears exclusively in the server-side recursion (Theorem 6.6), while both Γ_{mn}^t and Ψ_{mn}^t are used in the client-side update (Theorem 6.8). Making these matrices differentially private means that the propagated covariances become slightly biased or inflated due to the added noise. This leads to an overestimation of uncertainty, which preserves the structure of the recursion but may reduce the accuracy of Granger causality. Nonetheless, the estimates remain valid under perturbed inputs, and the user can control the privacy-utility trade-off via the (ϵ, δ) parameters.

D.5 COMPLEXITY ANALYSIS

Let each client m have state dimension p_m and data dimension d_m , with M clients in total. We first provide the computation and communication complexity per client.

(1) Computation. In a naive implementation, updating the full parameter covariance matrix $\Sigma_\theta^t \in \mathbb{R}^{p_m d_m \times p_m d_m}$ would incur $\mathcal{O}(p_m^2 d_m^2)$ computation per round due to matrix-matrix multiplications and Kronecker product evaluations. However, our implementation can avoid this cost by exploiting the structure of low-rank quantities such as $y^t y^{t\top}$ (which is rank-1), when $d_m \gg p_m$. As a result, matrix-vector products involving these terms can be computed without materializing the full $p_m d_m \times p_m d_m$ matrices. If we adopt a block-diagonal or factored representation of Σ_θ^t across layers or time steps, the per-client computational cost is further reduced to $\mathcal{O}(p_m d_m^2)$ per iteration.

(2) Communication. In terms of communication, our framework introduces no additional overhead compared to vanilla FedGC. Therefore, total communication cost per client remains $\mathcal{O}(p_m)$, matching that of FedGC. No additional communication is needed for tracking uncertainty.

All M clients. Across all M clients, the total additional computation scales as $\mathcal{O}\left(\sum_{m=1}^M p_m d_m^2\right)$ per round, while communication remains unchanged at $\mathcal{O}\left(\sum_{m=1}^M p_m\right)$. These properties make the method scalable to realistic federated settings.

E EXTENSIONS TO THE FRAMEWORK

E.1 NON-LINEAR MODELS

Our current theory is derived under the assumption of linear time-invariant dynamics and Gaussian noise, consistent with the original FedGC model. Nonetheless, we emphasize that the structure of our uncertainty propagation is not inherently tied to linearity; it generalizes to any dynamical model with well-defined latent state compression and state transition.

Bayesian Models. We consider an extension to *nonlinear but stationary* dynamical systems. Non-linear dynamics being highly model-specific, we first show that our theoretical structure carries over to two settings: (1) Extended Kalman Filters (EKF) - non-linear extension of the linear KF based formulation, and (2) Gaussian Processes (GPs) - for kernelized settings. Both EKF, and GP-based

frameworks support general nonlinear transition functions while admitting tractable recursive expressions analogous to Theorem 6.8.

(1) Extended Kalman Filter (EKF). EKF approximates nonlinear state transitions via first-order Taylor expansions. Consider the system:

$$h_c^t = f(h_c^{t-1}) + w^t, \quad y^t = g(h_c^t) + v^t, \quad w^t \sim \mathcal{N}(0, Q), \quad v^t \sim \mathcal{N}(0, R),$$

where $f : \mathbb{R}^p \rightarrow \mathbb{R}^p$, $g : \mathbb{R}^p \rightarrow \mathbb{R}^d$, with $d \gg p$, are smooth nonlinear functions. The prediction step linearizes f around the filtered mean:

$$\hat{h}_c^{t|t-1} = f(\hat{h}_c^{t-1}), \quad F^{t-1} := \left. \frac{\partial f}{\partial h} \right|_{h=\hat{h}_c^{t-1}}, \quad P^{t|t-1} = F^{t-1} P^{t-1} F^{t-1\top} + Q.$$

For the observation model, define $J_g := \left. \frac{\partial g}{\partial h} \right|_{h=\hat{h}_c^{t|t-1}}$. The update step proceeds via:

$$K^t = P^{t|t-1} J_g^\top (J_g P^{t|t-1} J_g^\top + R)^{-1}, \quad \hat{h}_c^t = \hat{h}_c^{t|t-1} + K^t (y^t - g(\hat{h}_c^{t|t-1})), \quad P^t = (I - K^t J_g) P^{t|t-1}.$$

We define the augmented representation used in FedGC:

$$\hat{h}_a^t = \hat{h}_c^t + \theta^t y^t, \quad \theta^t \in \mathbb{R}^{p \times d}, \quad v^t := \text{vec}(\theta^t).$$

Following our update rule, $v^{t+1} = v^t - \eta_1 \nabla_{v^t} \mathcal{L}_t^{\text{local}} - \eta_2 \nabla_{v^t} \mathcal{L}_t^{\text{server}}$, the parameter covariance update becomes:

$$\Sigma_\theta^{t+1} = H^t \Sigma_\theta^t H^{t\top} + G^t P^t G^{t\top} + H^t \Lambda^t G^{t\top} + G^t \Lambda^{t\top} H^t + P^t \Sigma_A^t P^{t\top},$$

where,

$$H^t = I_{pd} - 2\eta_1 (y^t y^{t\top}) \otimes (J_g^\top J_g) - 2\eta_2 (y^t y^{t\top}) \otimes (A_{mm}^\top A_{mm}),$$

$$G^t = 2\eta_1 (y^t \otimes J_g^\top),$$

$$P^t = \text{Cov}(\hat{h}_c^t),$$

$$\Lambda^t = \text{Cov}(v^t, \hat{h}_a^t).$$

This recursion is structurally identical to Theorem 6.8, with C and P replaced by the Jacobian J_g and EKF posterior P^t , respectively.

(2) Gaussian Process (GP). Consider a GP-based transition model, where $f \sim \mathcal{GP}(0, k(\cdot, \cdot))$ governs the latent dynamics:

$$h_c^t = f(h_c^{t-1}) + w^t.$$

Conditioning on past data $\{h_c^i, h_c^{i-1}\}_{i=1}^n$, the GP posterior yields:

$$h_c^t \sim \mathcal{N}(\mu_t^f, \Sigma_t^f + Q), \quad \mu_t^f = k^\top(h_c^{t-1}) K^{-1} \mathbf{h}, \quad \Sigma_t^f = k(h_c^{t-1}, h_c^{t-1}) - k^\top(h_c^{t-1}) K^{-1} k(h_c^{t-1}),$$

where K is the kernel matrix over $\{h_c^{i-1}\}$. We again define:

$$\hat{h}_a^t = \mu_t^f + \theta^t y^t.$$

The parameter covariance propagates as:

$$\Sigma_\theta^{t+1} = H^t \Sigma_\theta^t H^{t\top} + G^t \Sigma_t^f G^{t\top} + H^t \Lambda^t G^{t\top} + G^t \Lambda^{t\top} H^t + P^t \Sigma_A^t P^{t\top},$$

with $J_g := \partial g / \partial h|_{h=\mu_t^f}$ and other terms as above. This recursion mirrors the EKF case, with GP posterior variance Σ_t^f replacing P^t , and confirms that Theorem 6.8 applies structurally to kernelized approximations.

Non-Bayesian Models. While EKF and GP admit analytic propagation of posterior uncertainty, general deep models (e.g., RNNs, transformers) do not support closed-form covariance updates. Techniques such as Monte Carlo dropout, ensemble variance, or variational inference yield scalar uncertainties but do not enable recursive tracking. Thus, incorporating these models into our federated uncertainty framework would require entirely new derivations and likely sacrifice the tractability of the current analysis.

E.2 RELAXING STATIONARY ASSUMPTIONS

Our main theoretical results are derived under the classical assumption that each client observes weakly stationary data with time-invariant first- and second-order moments. This assumption enables closed-form uncertainty propagation through the FedGC recursions. Here we clarify how the framework behaves when stationarity is mildly violated and how it can be adapted to non-stationary settings.

Exponentially weighted moments. The key observation is that all uncertainty recursions remain valid when the empirical moments are replaced by *exponentially weighted* (EWMA) moments,

$$\mu_t = (1 - \lambda)\mu_{t-1} + \lambda x_t, \quad \Sigma_t = (1 - \lambda)\Sigma_{t-1} + \lambda x_t x_t^\top,$$

for a forgetting factor $0 < \lambda < 1$. Suppose the underlying data possess slowly drifting moments (μ_t^*, Σ_t^*) with bounded temporal variation $\|\mu_t^* - \mu_{t-1}^*\| \leq \delta$ and $\|\Sigma_t^* - \Sigma_{t-1}^*\| \leq \delta$. Standard results for stochastic approximation imply

$$\|\mu_t - \mu_t^*\| = O(\delta/\lambda), \quad \|\Sigma_t - \Sigma_t^*\| = O(\delta/\lambda).$$

Thus, EWMA moments track the true time-varying moments whenever the drift is slower than the forgetting rate. Replacing stationary moments by EWMA moments preserves the algebraic form of our uncertainty recursions; the same propagation equations apply with (μ_t, Σ_t) in place of fixed moments.

Beyond EWMA, other forms of non-stationarity can be incorporated by modifying the moment-estimation step. Examples include sliding-window estimators, seasonally adjusted or periodic-window estimators, trend-filtered or total-variation-regularized moment updates, and online convex-combination estimators for abrupt regime changes. Neural-network-based models can also accommodate complex non-stationarity, but doing so would require deriving a new set of uncertainty-propagation equations beyond the linear FedGC framework. A systematic treatment of these extensions is an important direction for future work.

F PROOFS

F.1 PROPOSITION 6.1

• **Proposition (Client Model-Client Data Dependence)** Assume $\text{Var}(y_m^{t-1}) > 0$. Then under the federated Granger-causality updates, $\Omega_m^t := \text{Cov}(v_m^t, y_m^t) \neq 0$.

Proof. From the gradient-descent update we have,

$$\begin{aligned} \theta_m^t &= \theta_m^{t-1} + 2\eta_1 (C_{mm} A_{mm})^\top (y_m^t - C_{mm} A_{mm} [\hat{h}_{m,c}^{t-1} + \theta_m^{t-1} y_m^{t-1}]) y_m^{t-1\top} \\ &\quad - 2\eta_2 A_{mm}^\top \left(A_{mm} [\hat{h}_{m,a}^{t-1} - \hat{h}_{m,c}^{t-1}] - \sum_{n \neq m} \hat{A}_{mn}^{t-1} \hat{h}_{n,c}^{t-1} \right) y_m^{t-1\top}. \end{aligned}$$

Rearrange to isolate dependence on y_m^{t-1} :

$$\theta_m^t = M + B y_m^{t-1\top},$$

where

$$\begin{aligned} M &= \theta_m^{t-1} + 2\eta_1 (C_{mm} A_{mm})^\top (y_m^t - C_{mm} A_{mm} \hat{h}_{m,c}^{t-1}) y_m^{t-1\top} - 2\eta_2 A_{mm}^\top \left(A_{mm} [\hat{h}_{m,a}^{t-1} - \hat{h}_{m,c}^{t-1}] \right) y_m^{t-1\top}, \\ B &= -2\eta_1 (C_{mm} A_{mm})^\top C_{mm} A_{mm} \theta_m^{t-1} - 2\eta_2 A_{mm}^\top \sum_{n \neq m} \hat{A}_{mn}^{t-1} \hat{h}_{n,c}^{t-1}. \end{aligned}$$

Thus $\theta_m^t = M + B y_m^{t-1\top}$ is an affine function of y_m^{t-1} .

Furthermore, the LTI measurement model gives

$$y_m^t = C_{mm} A_{mm} \hat{h}_{m,a}^{t-1} + \text{Var} \epsilon_m^t = C_{mm} A_{mm} (\hat{h}_{m,c}^{t-1} + \theta_m^{t-1} y_m^{t-1}) + \text{Var} \epsilon_m^t.$$

Rearranging we obtain,

$$y_m^t = N + D y_m^{t-1} + \text{Var } \epsilon_m^t,$$

where,

$$N = C_{mm} A_{mm} \hat{h}_{m,c}^{t-1}, \quad D = C_{mm} A_{mm} \theta_m^{t-1}.$$

Thus y_m^t is also an affine function of y_m^{t-1} .

Let $u = y_m^{t-1}$. Then

$$\theta_m^t = M + B u^\top, \quad y_m^t = N + D u + \text{Var } \epsilon_m^t.$$

Since $\text{Var } \epsilon_m^t$ is zero-mean and independent of u , we have

$$\text{Cov}(\theta_m^t, y_m^t) = \text{Cov}(M + B u^\top, N + D u) = B \text{Cov}(u^\top, u) D^\top = B \text{Var}(u) D^\top.$$

By assumption $\text{Var}(u) = \text{Var}(y_m^{t-1}) > 0$, and B, D are nonzero (since the update and measurement matrices are full-rank). Therefore $\text{Cov}(\theta_m^t, y_m^t) = B \text{Var}(y_m^{t-1}) D^\top \neq 0$.

Hence $\Omega_m^t \neq 0$, as claimed. \square

F.2 PROPOSITION 6.2

• **Proposition (Client Model-Client State Dependence)** *Let $\Lambda_m^t = \text{Cov}(v_m^t, \hat{h}_{m,a}^t)$. Then we have the following recursion within the client, $\Lambda_m^t = \Sigma_{\theta_m}^t (I_{d_m} \otimes \mu_{y_m}^t) + \Omega_m^t (\mu_{v_m}^t \otimes I_{d_m})$,*

Proof. We prove the recursion for $\Lambda_m^t = \text{Cov}(v_m^t, \hat{h}_{m,a}^t)$ using the paper's definitions:

From Table 1 and Eq. (4), we have the following definitions,

$$\begin{aligned} \Sigma_{\theta_m}^t &:= \text{Var}(v_m^t) = \mathbb{E}[v_m^t v_m^{t\top}] - \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} \\ \Omega_m^t &:= \text{Cov}(v_m^t, y_m^t) = \mathbb{E}[v_m^t y_m^{t\top}] - \mu_{\theta_m}^t \mu_{y_m}^{t\top} \\ \hat{h}_{m,a}^t &:= \hat{h}_{m,c}^t + \theta_m^t y_m^t = \hat{h}_{m,c}^t + (y_m^{t\top} \otimes I_{p_m}) v_m^t \end{aligned}$$

Expanding Λ_m^t from first principles, and using the definition of $\hat{h}_{m,a}^t$ we have,

$$\begin{aligned} \Lambda_m^t &= \mathbb{E}[v_m^t \hat{h}_{m,a}^{t\top}] - \mu_{\theta_m}^t \mu_{\hat{h}_{m,a}}^{t\top} \\ &= \mathbb{E}\left[v_m^t \left(\hat{h}_{m,c}^{t\top} + v_m^{t\top} (y_m^t \otimes I_{p_m})\right)\right] - \mu_{\theta_m}^t \left(\hat{h}_{m,c}^{t\top} + \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m})\right) \\ &= \mathbb{E}[v_m^t v_m^{t\top} (y_m^t \otimes I_{p_m})] - \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) \end{aligned}$$

Analyzing the key expectation terms and substituting the definitions of $\Sigma_{\theta_m}^t$, and Ω_m^t we have,

$$\begin{aligned} \mathbb{E}[v_m^t v_m^{t\top} (y_m^t \otimes I_{p_m})] &= \mathbb{E}\left[(\Sigma_{\theta_m}^t + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}) (y_m^t \otimes I_{p_m})\right] \quad (\text{since } \mathbb{E}[v_m^t v_m^{t\top}] = \Sigma_{\theta_m}^t + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}) \\ &= \Sigma_{\theta_m}^t (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) \\ &\quad + \mathbb{E}[(v_m^t - \mu_{\theta_m}^t)(v_m^t - \mu_{\theta_m}^t)^\top (y_m^t - \mu_{y_m}^t \otimes I_{p_m})] \\ &= \Sigma_{\theta_m}^t (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) \\ &\quad + \mu_{\theta_m}^t \mathbb{E}[(v_m^t - \mu_{\theta_m}^t)(y_m^t - \mu_{y_m}^t)^\top] \otimes I_{p_m} \\ &= \Sigma_{\theta_m}^t (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \Omega_m^{t\top} \otimes I_{p_m} \end{aligned}$$

Substituting the expression for $\mathbb{E}[v_m^t v_m^{t\top} (y_m^t \otimes I_{p_m})]$ (obtained above), back into Λ_m^t we have,

$$\begin{aligned} \Lambda_m^t &= (\Sigma_{\theta_m}^t (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \Omega_m^{t\top} \otimes I_{p_m}) \\ &\quad - \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} (\mu_{y_m}^t \otimes I_{p_m}) \\ &= \Sigma_{\theta_m}^t (\mu_{y_m}^t \otimes I_{p_m}) + \mu_{\theta_m}^t \Omega_m^{t\top} \otimes I_{p_m} \end{aligned}$$

Recognizing that $\Omega_m^{t\top} \otimes I_{p_m} = (\Omega_m^t \otimes I_{p_m})^\top$, and simplifying the second term we obtain,

$$\mu_{\theta_m}^t \Omega_m^{t\top} \otimes I_{p_m} = (\Omega_m^t \otimes I_{p_m}) \mu_{\theta_m}^t = \Omega_m^t (\mu_{\theta_m}^t \otimes I_{d_m})$$

Substituting this into the expression for Λ_m^t we have,

$$\Lambda_m^t = \Sigma_{\theta_m}^t (I_{d_m} \otimes \mu_{y_m}^t) + \Omega_m^t (\mu_{\theta_m}^t \otimes I_{d_m})$$

□

F.3 LEMMA 6.3

Assumption (A6): The client’s augmented hidden state changes only by a small amount between two consecutive time-steps $\Delta h_m^t := \hat{h}_{m,a}^{t+1} - \hat{h}_{m,a}^t$ satisfying $\|\Delta h_m^t\|_2 \leq \varepsilon$ with ε is small.

• **Lemma (Client State-Sever Model Dependence)** *The cross-covariance term $\Gamma_{mn}^t := \text{Cov}(a_{mn}^t, \hat{h}_{m,a}^t)$ follows the recursive equation: $\Gamma_{mn}^{t+1} = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t$ where, $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$, $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$, and $\Sigma_{h_m}^t = \text{Var}(\hat{h}_{m,a}^t)$.*

Proof. Gradient descent on the quadratic loss L_s with step size γ gives,

$$a_{mn}^{t+1} = D_n^t a_{mn}^t + 2\gamma B_{mn}^t \hat{h}_{m,a}^t,$$

where $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$ and $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$.

Taking the column covariance with $\hat{h}_{m,a}^t$ yields the shifted covariance term

$$\tilde{\Gamma}_{mn}^{t+1} := \text{Cov}(a_{mn}^{t+1}, \hat{h}_{m,a}^t) = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t.$$

By Assumption (A6),

$$\hat{h}_{m,a}^{t+1} = \hat{h}_{m,a}^t + \Delta h_m^t, \quad \|\Delta h_m^t\|_2 \leq \varepsilon.$$

Defining, $\Gamma_{mn}^{t+1} := \text{Cov}(a_{mn}^{t+1}, \hat{h}_{m,a}^{t+1})$ and expanding, we get,

$$\Gamma_{mn}^{t+1} = \text{Cov}(a_{mn}^{t+1}, \hat{h}_{m,a}^t + \Delta h_m^t)$$

Let us define $E_{mn}^t := \text{Cov}(a_{mn}^{t+1}, \hat{h}_{m,a}^t + \Delta h_m^t)$. Thne the matrix Cauchy–Schwarz inequality gives,

$$\|E_{mn}^t\|_2 \leq \sqrt{\text{tr}(\Sigma_{A_{mn}}^{t+1})} \varepsilon = O(\varepsilon).$$

Using the above expression we obtain,

$$\Gamma_{mn}^{t+1} = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t + O(\varepsilon),$$

If ε is small (Assumption (A6)), we get,

$$\Gamma_{mn}^{t+1} = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t$$

□

F.4 LEMMA 6.4

• **Lemma (Client Model-Server Model Dependence)** *The term $\Psi_{mn}^t := \text{Cov}(a_{mn}^t, v_m^t)$ evolves as, $\Psi_{mn}^{t+1} = D_n^t \Psi_{mn}^t H_m^{t\top} + D_n^t \Gamma_{mn}^t G_m^{t\top} - D_n^t \Sigma_{A_{mn}}^t P_m^{t\top} + 2\gamma B_{mn}^t \Lambda_m^t H_m^{t\top} + 2\gamma B_{mn}^t \Sigma_{h_m}^t G_m^{t\top} - 2\gamma B_{mn}^t \Gamma_{mn}^{t\top} P_m^{t\top}$, with the following gain matrices, $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$,*

$$D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I, G_m^t = 2\eta_1 (y_m^t \otimes (C_{mm} A_{mm})^\top), P_m^t = -2\eta_2 (y_m^t \otimes A_{mm}^\top)$$

$$\text{and } H_m^t = I_{p_m d_m} - 2\eta_1 (y_m^t y_m^{t\top}) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (y_m^t y_m^{t\top}) \otimes (A_{mm}^\top A_{mm})$$

Proof. From the loss L_s one gradient–descent step with stepsize γ gives,

$$a_{mn}^{t+1} = D_n^t a_{mn}^t + 2\gamma B_{mn}^t \hat{h}_{m,a}^t.$$

The update $\theta_m^{t+1} = \theta_m^t - \eta_1 \nabla_{\theta_m} (L_m)_a - \eta_2 \nabla_{\theta_m} L_s$ is linear in $(\theta_m^t, \hat{h}_{m,a}^t, a_{mn}^t)$; in vectorized form,

$$v_m^{t+1} = H_m^t v_m^t + G_m^t \hat{h}_{m,a}^t - P_m^t a_{mn}^t.$$

Compute $\Psi_{mn}^{t+1} = \text{Cov}(a_{mn}^{t+1}, v_m^{t+1})$ using the above two equations,

$$\begin{aligned} \Psi_{mn}^{t+1} &= \text{Cov}(D_n^t a_{mn}^t + 2\gamma B_{mn}^t \hat{h}_{m,a}^t, H_m^t v_m^t + G_m^t \hat{h}_{m,a}^t - P_m^t a_{mn}^t) \\ &= D_n^t \text{Cov}(a_{mn}^t, v_m^t) H_m^{t\top} + D_n^t \text{Cov}(a_{mn}^t, \hat{h}_{m,a}^t) G_m^{t\top} - D_n^t \text{Cov}(a_{mn}^t, a_{mn}^t) P_m^{t\top} \\ &\quad + 2\gamma B_{mn}^t \text{Cov}(\hat{h}_{m,a}^t, v_m^t) H_m^{t\top} + 2\gamma B_{mn}^t \text{Cov}(\hat{h}_{m,a}^t, \hat{h}_{m,a}^t) G_m^{t\top} - 2\gamma B_{mn}^t \text{Cov}(\hat{h}_{m,a}^t, a_{mn}^t) P_m^{t\top} \\ &= D_n^t \Psi_{mn}^t H_m^{t\top} + D_n^t \Gamma_{mn}^t G_m^{t\top} - D_n^t \Sigma_{A_{mn}}^t P_m^{t\top} + 2\gamma B_{mn}^t \Lambda_m^t H_m^{t\top} + 2\gamma B_{mn}^t \Sigma_{h_m}^t G_m^{t\top} - 2\gamma B_{mn}^t \Gamma_{mn}^t P_m^{t\top}. \end{aligned}$$

Grouping the six contributions yields the given recursion of Ψ_{mn}^{t+1} . \square

F.5 LEMMA 6.5

• **Lemma (Uncertainty in Client to Server Communication)** Let $\kappa_m = \text{tr}(\Sigma_{y_m}^t) + \|\mu_{y_m}^t\|^2$. Then the variance in the $\hat{h}_{m,a}^t$ is given by, $\Sigma_{h_m}^t = \kappa_m \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}$.

Proof. From the definition of augmented client states we have, $q_m^t := \hat{h}_{m,a}^t - \hat{h}_{m,c}^t = \theta_m^t y_m^t$, which in vectorized form is $q_m^t := Y_m^t v_m^t$, with $Y_m^t := (y_m^{t\top} \otimes I_{p_m})$. Therefore, by definition we have,

$$\text{Var}(q_m^t) = \mathbb{E}[Y_m^t v_m^t v_m^{t\top} Y_m^{t\top}] - \mathbb{E}[Y_m^t v_m^t] \mathbb{E}[Y_m^t v_m^t]^\top.$$

Computing the second moment we have,

$$\begin{aligned} \mathbb{E}[Y_m^t v_m^t v_m^{t\top} Y_m^{t\top}] &= \mathbb{E}[(y_m^t \otimes I_{p_m}) v_m^t v_m^{t\top} (y_m^{t\top} \otimes I_{p_m})] \\ &= \mathbb{E}[(y_m^t y_m^{t\top}) \otimes (v_m^t v_m^{t\top})] \\ &= \mathbb{E}[y_m^t y_m^{t\top}] \otimes \mathbb{E}[v_m^t v_m^{t\top}] + \text{Cov}(y_m^t \otimes v_m^t) \\ &= (\Sigma_{y_m}^t + \mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes (\Sigma_{\theta_m}^t + \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}) \\ &\quad + \Omega_m^t \otimes (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \otimes \Omega_m^{t\top}. \end{aligned}$$

The first moment is given by,

$$\mathbb{E}[Y_m^t v_m^t] = (\mu_{y_m}^t \otimes I_{p_m}) \mu_{\theta_m}^t + \Omega_m^t$$

Computing the outer product of the first moments we have,

$$\begin{aligned} \mathbb{E}[Y_m^t v_m^t] \mathbb{E}[Y_m^t v_m^t]^\top &= (\mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes (\mu_{\theta_m}^t \mu_{\theta_m}^{t\top}) \\ &\quad + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}. \end{aligned}$$

Subtracting outer product of first moment from second moment we obtain,

$$\begin{aligned} \text{Var}(q_m^t) &= [(\Sigma_{y_m}^t + \mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes \Sigma_{\theta_m}^t] + [(\Sigma_{y_m}^t + \mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}] \\ &\quad - (\mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes (\mu_{\theta_m}^t \mu_{\theta_m}^{t\top}) \\ &\quad + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}. \end{aligned}$$

The term $(\Sigma_{y_m}^t + \mu_{y_m}^t \mu_{y_m}^{t\top}) \otimes \mu_{\theta_m}^t \mu_{\theta_m}^{t\top} - \mu_{y_m}^t \mu_{y_m}^{t\top} \otimes \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}$ simplifies to $\Sigma_{y_m}^t \otimes \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}$. By design, $\Sigma_{y_m}^t \otimes \mu_{\theta_m}^t \mu_{\theta_m}^{t\top}$ is absorbed into $\kappa_m \Sigma_{\theta_m}^t$ via trace normalization, leaving,

$$\text{Var}(q_m^t) = \kappa_m \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}.$$

Using the definition of q_m^t we have,

$$\text{Var}(\hat{h}_{m,a}^t - \hat{h}_{m,c}^t) = \kappa_m \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}.$$

Since $\hat{h}_{m,c}^t$ is deterministic, we have,

$$\Sigma_{h_m}^t := \text{Var}(\hat{h}_{m,a}^t) = \kappa_m \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t\top}.$$

\square

F.6 LEMMA 6.6

• **Lemma (Uncertainty in Server to Client Communication)** *With notation as above, the uncertainty in the gradient communicated by the server is given by, $\text{Var}(g_{m,s}^{t+1}) = A_{mm}^\top U^t A_{mm}$, where $U^t = A_{mm} \Sigma_{h_m}^t A_{mm}^\top + \sum_{n \neq m} (h_{n,c}^t h_{n,c}^{t\top}) \Sigma_{A_{mn}}^t - 2 \sum_{n \neq m} A_{mm} \Gamma_{mn}^t h_{n,c}^{t\top}$.*

Proof. Let $r^t := A_{mm}(\hat{h}_{m,a}^t - \hat{h}_{m,c}^t) - \sum_{n \neq m} \hat{A}_{mn}^t \hat{h}_{n,c}^t$. We know that, $g_{m,s}^{t+1} = A_{mm}^\top r^t$. Then

$$\text{Var}(g_{m,s}^{t+1}) = A_{mm}^\top \text{Var}(r^t) A_{mm}.$$

We compute

$$\text{Var}(r^t) := \text{Var}(A_{mm} \hat{h}_{m,a}^t) + \sum_{n \neq m} \text{Var}(\hat{A}_{mn}^t \hat{h}_{n,c}^t) - 2 \sum_{n \neq m} \text{Cov}(A_{mm} \hat{h}_{m,a}^t, \hat{A}_{mn}^t \hat{h}_{n,c}^t),$$

Since $\hat{h}_{m,c}^t$ is deterministic. We have,

$$\text{Var}(A_{mm} \hat{h}_{m,a}^t) = A_{mm} \Sigma_{h_m}^t A_{mm}^\top,$$

$$\text{Var}(\hat{A}_{mn}^t \hat{h}_{n,c}^t) = (h_{n,c}^t h_{n,c}^{t\top}) \Sigma_{A_{mn}}^t; \text{Cov}(A_{mm} \hat{h}_{m,a}^t, \hat{A}_{mn}^t \hat{h}_{n,c}^t) = A_{mm} \Gamma_{mn}^t h_{n,c}^{t\top}.$$

Putting these into $\text{Var}(r^t)$ gives exactly

$$U^t = A_{mm} \Sigma_{h_m}^t A_{mm}^\top + \sum_{n \neq m} (h_{n,c}^t h_{n,c}^{t\top}) \Sigma_{A_{mn}}^t - 2 \sum_{n \neq m} A_{mm} \Gamma_{mn}^t h_{n,c}^{t\top}.$$

Hence $\text{Var}(g_{m,s}^{t+1}) = A_{mm}^\top U^t A_{mm}$, as claimed. \square

F.7 THEOREM 6.7

• **Theorem (Uncertainty Propagation within the Server)** *The server model parameter a_{mn}^t 's covariance $\Sigma_{A_{mn}}^t$ evolves as, $\Sigma_{A_{mn}}^{t+1} = D_n^t \Sigma_{A_{mn}}^t D_n^{t\top} + 4\gamma^2 (\hat{h}_{n,c}^t \otimes A_{mm}) \Sigma_{h_m}^t (\hat{h}_{n,c}^t \otimes A_{mm})^\top + 2\gamma (D_n^t \Gamma_{mn}^t B_{mn}^{t\top} + B_{mn}^t \Gamma_{mn}^{t\top} D_n^{t\top})$ with $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$, and $B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}$*

Proof. At round t , the server gradient update as follows:

$$\hat{A}_{mn}^{t+1} = \hat{A}_{mn}^t (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) + 2\gamma A_{mm} [\hat{h}_{m,a}^t - \hat{h}_{m,c}^t] \hat{h}_{n,c}^{t\top} - 2\gamma \sum_{p \neq m, n} \hat{A}_{mp}^t \hat{h}_{p,c}^t \hat{h}_{n,c}^{t\top}.$$

Under Assumption (A1) (off-diagonal blocks independent), the last summation term contributes no covariance with \hat{A}_{mn}^t and can be omitted when computing $\text{Var}(\hat{A}_{mn}^{t+1})$.

Apply $\text{Vec}(\cdot)$ and use the property that: $\text{Vec}(XB) = (B^\top \otimes I) \text{Vec}(X)$ and $\text{Vec}(AX) = (I \otimes A) \text{Vec}(X)$ for any three matrices A, B, X .

We obtain the following after vectorization,

$$\text{Vec}(\hat{A}_{mn}^{t+1}) = \left((I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I \right) \text{Vec}(\hat{A}_{mn}^t) + 2\gamma (\hat{h}_{n,c}^t \otimes A_{mm}) \hat{h}_{m,a}^t.$$

We then define,

$$D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I, \quad B_{mn}^t = \hat{h}_{n,c}^t \otimes A_{mm}, \quad a_{mn}^t = \text{Vec}(\hat{A}_{mn}^t).$$

Then the vectorized update is given by,

$$a_{mn}^{t+1} = D_n^t a_{mn}^t + 2\gamma B_{mn}^t \hat{h}_{m,a}^t.$$

We wish to compute $\Sigma_{A_{mn}}^{t+1} = \text{Var}(a_{mn}^{t+1})$.

Using the property $\text{Var}[X + Z] = \text{Var}[X] + \text{Var}[Z] + \text{Cov}(X, Z) + \text{Cov}(Z, X)$ for any two vectors X , and Z .

We set, $X = D_n^t a_{mn}^t$, $Z = 2\gamma B_{mn}^t \hat{h}_{m,a}^t$. Then,
 $\Sigma_{A_{mn}}^{t+1} = \text{Var}[X] + \text{Var}[Z] + \text{Cov}(X, Z) + \text{Cov}(Z, X)$.

(1) *Variance of X* : D_n^t is deterministic, so

$$\text{Var}[X] = D_n^t \text{Var}(a_{mn}^t) D_n^{t\top} = D_n^t \Sigma_{A_{mn}}^t D_n^{t\top}.$$

(2) *Variance of Z* : $\hat{h}_{n,c}^t$ and A_{mm} are fixed at round t , hence

$$\text{Var}[Z] = 4\gamma^2 B_{mn}^t \text{Var}(\hat{h}_{m,a}^t) B_{mn}^{t\top} = 4\gamma^2 (\hat{h}_{n,c}^t \otimes A_{mm}) \Sigma_{h_m}^t (\hat{h}_{n,c}^t \otimes A_{mm})^\top.$$

(3) *Cross-covariance terms*: Since D_n^t and B_{mn}^t are deterministic,

$$\text{Cov}(X, Z) = 2\gamma D_n^t \text{Cov}(a_{mn}^t, \hat{h}_{m,a}^t) B_{mn}^{t\top} = 2\gamma D_n^t \Gamma_{mn}^t B_{mn}^{t\top},$$

$$\text{Cov}(Z, X) = 2\gamma B_{mn}^t \text{Cov}(\hat{h}_{m,a}^t, a_{mn}^t) D_n^{t\top} = 2\gamma B_{mn}^t \Gamma_{mn}^{t\top} D_n^{t\top}.$$

Adding the four contributions we obtain,

$$\begin{aligned} \Sigma_{A_{mn}}^{t+1} &= D_n^t \Sigma_{A_{mn}}^t D_n^{t\top} + 4\gamma^2 (\hat{h}_{n,c}^t \otimes A_{mm}) \Sigma_{h_m}^t (\hat{h}_{n,c}^t \otimes A_{mm})^\top \\ &\quad + 2\gamma (D_n^t \Gamma_{mn}^t B_{mn}^{t\top} + B_{mn}^t \Gamma_{mn}^{t\top} D_n^{t\top}) \end{aligned}$$

□

F.8 THEOREM 6.8

• **Theorem (Uncertainty Propagation within the Client)** *The client-parameter covariance $\Sigma_{\theta_m}^t$ obeys the following recursion: $\Sigma_{\theta_m}^t = H_m^{t-1} \Sigma_{\theta_m}^{t-1} H_m^{t-1\top} + G_m^{t-1} \Sigma_{h_m}^{t-1} G_m^{t-1\top} + (X_m + X_m^\top) - \sum_{n \neq m} (Y_{mn} + Y_{mn}^\top) - \sum_{n \neq m} (Z_{mn} + Z_{mn}^\top) + \sum_{n \neq m} P_m^{t-1} \Sigma_{A_{mn}}^{t-1} P_m^{t-1\top}$,*

where, $X_m = H_m^{t-1} \Lambda_m^{t-1} G_m^{t-1\top}$, $Y_{mn} = H_m^{t-1} \Psi_{mn}^{t-1} P_m^{t-1\top}$, $Z_{mn} = G_m^{t-1} \Gamma_{mn}^{t-1} P_m^{t-1\top}$,
and, $G_m^{t-1} := 2\eta_1 (y_m^{t-1} \otimes (C_{mm} A_{mm})^\top)$, $P_m^{t-1} := -2\eta_2 (y_m^{t-1} \otimes A_{mm}^\top)$, $H_m^{t-1} := I_{p_m d_m} - 2\eta_1 (y_m^{t-1} y_m^{t-1\top}) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (y_m^{t-1} y_m^{t-1\top}) \otimes (A_{mm}^\top A_{mm})$

Proof. All random variables, distributional assumptions (A1–A6) and second-moment symbols $\Sigma_{\theta_m}^t, \Sigma_{h_m}^t, \Sigma_{A_{mn}}^t, \Lambda_m^t, \Psi_{mn}^t, \Gamma_{mn}^t, \Omega_m^t$ are defined in Section 5.

Computing the analytical values of $\nabla_{\theta_m}^t (L_m)_a$, and $\nabla_{\theta_m}^t L_s$, and substituting them in Eq (5) of Section 3.1, we have the following client model update (for the FedGC framework):

$$\theta_m^t = \theta_m^{t-1} + 2\eta_1 (C_{mm} A_{mm})^\top \left(y_m^{t-1} - C_{mm} A_{mm} \hat{h}_{m,c}^{t-1} \right) y_m^{t-1\top} - 2\eta_2 A_{mm}^\top \left(\sum_{n \neq m} A_{mn}^{t-1} \hat{h}_{n,c}^{t-1} \right) y_m^{t-1\top}.$$

Let $v_m^t := \text{Vec}(\theta_m^t)$. Using the property $\text{Vec}(AXB) = (B^\top \otimes A) \text{Vec}(X)$ and $\text{Vec}(A \hat{h}) = (\hat{h}^\top \otimes I) \text{Vec}(A)$ for any matrices A, B, X and vector \hat{h} , we obtain the following:

$$v_m^t = H_m^{t-1} v_m^{t-1} + G_m^{t-1} \hat{h}_{m,a}^{t-1} + P_m^{t-1} \sum_{n \neq m} a_{mn}^{t-1} + u_m^{t-1},$$

where these matrices (H_m^t, G_m^t, P_m^t) are exactly those stated in the theorem. We define $u_m^{t-1} := 2\eta_1 (y_m^{t-1} \otimes (C_{mm} A_{mm})^\top) (y_m^{t-1} - C_{mm} A_{mm} \hat{h}_{m,c}^{t-1})$. By assumption (A4) u_m^{t-1} has mean 0 and vanishing covariance: $\mathbb{E}[u_m^{t-1}] = 0$, $\text{Var}(u_m^{t-1}) = 0$.

First, we have the following:

$$\text{Var}(H_m^{t-1} v_m^{t-1}) = H_m^{t-1} \Sigma_{\theta_m}^{t-1} H_m^{t-1\top}, \quad \text{Var}(G_m^{t-1} \hat{h}_{m,a}^{t-1}) = G_m^{t-1} \Sigma_{h_m}^{t-1} G_m^{t-1\top}.$$

We also denote $S_m^{t-1} := \sum_{n \neq m} a_{mn}^{t-1}$. Independence of different off-diagonal blocks a_{mn}^{t-1} in assumption **(A2)** yields $\text{Var}(S_m^{t-1}) = \sum_{n \neq m} \Sigma_{A_{mn}}^{t-1}$; hence

$$\text{Var}(P_m^{t-1} S_m^{t-1}) = \sum_{n \neq m} P_m^{t-1} \Sigma_{A_{mn}}^{t-1} P_m^{t-1\top}.$$

Independence assumptions **(A2)** imply that cross terms with different client indices cancel. The only non-zero covariances are

$$\begin{aligned} \text{Cov}(Hv, G\hat{h}) &= H_m^{t-1} \Lambda_m^{t-1} G_m^{t-1\top} = X_m, \\ \text{Cov}(Hv, PS) &= \sum_{n \neq m} H_m^{t-1} \Psi_{mn}^{t-1} P_m^{t-1\top} = \sum_{n \neq m} Y_{mn}, \\ \text{Cov}(G\hat{h}, PS) &= \sum_{n \neq m} G_m^{t-1} \Gamma_{mn}^{t-1} P_m^{t-1\top} = \sum_{n \neq m} Z_{mn}. \end{aligned}$$

Each term X_m, Y_{mn}, Z_{mn} appears together with its transpose in the variance expansion. Applying $\text{Var}(\cdot)$ to $\text{Vec}(\theta_m)$, and using $\text{Var}(u_m^{t-1}) = 0$, we obtain:

$$\begin{aligned} \Sigma_{\theta_m}^t &= H_m^{t-1} \Sigma_{\theta_m}^{t-1} H_m^{t-1\top} + G_m^{t-1} \Sigma_{h_m}^{t-1} G_m^{t-1\top} + (X_m + X_m^\top) \\ &\quad - \sum_{n \neq m} (Y_{mn} + Y_{mn}^\top) - \sum_{n \neq m} (Z_{mn} + Z_{mn}^\top) + \sum_{n \neq m} P_m^{t-1} \Sigma_{A_{mn}}^{t-1} P_m^{t-1\top}. \end{aligned}$$

□

F.9 PROPOSITION 7.1

• Proposition (Gain Matrices Convergence) *Under the above assumptions, the gain matrices used in Section 6 converges as, $\lim_{t \rightarrow \infty} (D_n^t, H_m^t, G_m^t, P_m^t) = (D_n, H_m, G_m, P_m)$ where, $D_n = (I - 2\gamma \hat{h}_{n,c} \hat{h}_{n,c}^\top) \otimes I$, $G_m = 2\eta_1 (\mu_{y_m} \otimes (C_{mm} A_{mm})^\top)$, $P_m = -2\eta_2 (\mu_{y_m} \otimes A_{mm}^\top)$ and $H_m = I_{p_m d_m} - 2\eta_1 (\mu_{y_m} \mu_{y_m}^\top) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) - 2\eta_2 (\mu_{y_m} \mu_{y_m}^\top) \otimes (A_{mm}^\top A_{mm})$.*

Proof. We prove the convergence of the gain matrices under the assumptions (provided in Section 7):

- (I) $\lim_{t \rightarrow \infty} y_m^t = \mu_{y_m}$ (client data converges)
- (II) $\lim_{t \rightarrow \infty} \hat{h}_{m,c}^t = \hat{h}_{m,c}$ (client state estimates converge)

First, for $D_n^t = (I - 2\gamma \hat{h}_{n,c}^t \hat{h}_{n,c}^{t\top}) \otimes I$ we have,

$$\begin{aligned} \lim_{t \rightarrow \infty} D_n^t &= \left(I - 2\gamma \left(\lim_{t \rightarrow \infty} \hat{h}_{n,c}^t \right) \left(\lim_{t \rightarrow \infty} \hat{h}_{n,c}^t \right)^\top \right) \otimes I \\ &= (I - 2\gamma \hat{h}_{n,c} \hat{h}_{n,c}^\top) \otimes I =: D_n \end{aligned}$$

Next for $G_m^t = 2\eta_1 (y_m^t \otimes (C_{mm} A_{mm})^\top)$ we have,

$$\begin{aligned} \lim_{t \rightarrow \infty} G_m^t &= 2\eta_1 \left(\left(\lim_{t \rightarrow \infty} y_m^t \right) \otimes (C_{mm} A_{mm})^\top \right) \\ &= 2\eta_1 (\mu_{y_m} \otimes (C_{mm} A_{mm})^\top) =: G_m \end{aligned}$$

Similarly for $P_m^t = -2\eta_2 (y_m^t \otimes A_{mm}^\top)$ we have,

$$\begin{aligned} \lim_{t \rightarrow \infty} P_m^t &= -2\eta_2 \left(\left(\lim_{t \rightarrow \infty} y_m^t \right) \otimes A_{mm}^\top \right) \\ &= -2\eta_2 (\mu_{y_m} \otimes A_{mm}^\top) =: P_m \end{aligned}$$

1512 Finally for H_m^t we have,
1513

$$1514 \lim_{t \rightarrow \infty} H_m^t = I_{p_m d_m} - 2\eta_1 \left(\left(\lim_{t \rightarrow \infty} y_m^t y_m^{t\top} \right) \otimes ((C_{mm} A_{mm})^\top C_{mm} A_{mm}) \right) \\ 1515 - 2\eta_2 \left(\left(\lim_{t \rightarrow \infty} y_m^t y_m^{t\top} \right) \otimes (A_{mm}^\top A_{mm}) \right) \\ 1516 \\ 1517$$

1518 Using the fact that $\lim_{t \rightarrow \infty} y_m^t y_m^{t\top} = \mu_{y_m} \mu_{y_m}^\top + \Sigma_{y_m}$ (from the stationary distribution), but under
1519 Assumption (A4) that Σ_{y_m} is constant, we get,
1520

$$1521 H_m = I_{p_m d_m} - 2\eta_1 (\mu_{y_m} \mu_{y_m}^\top \otimes (C_{mm} A_{mm})^\top C_{mm} A_{mm}) \\ 1522 - 2\eta_2 (\mu_{y_m} \mu_{y_m}^\top \otimes A_{mm}^\top A_{mm}) \\ 1523$$

□

1526 F.10 PROPOSITION 7.2

1527 • **Proposition** If $\rho(D_n) < 1$ and $\rho(H_m) < 1$ then we have, $\lim_{t \rightarrow \infty} (\Gamma_{mn}^t, \Psi_{mn}^t) = (\Gamma_{mn}^\infty, \Psi_{mn}^\infty)$
1528 with $\Gamma_{mn}^\infty = (I - D_n)^{-1} 2\gamma B_{mn} \Sigma_{h_m}^\infty$, & $\Psi_{mn}^\infty = (I - H_m \otimes D_n)^{-1} \text{Vec}(D_n \Gamma_{mn}^\infty G_m^\top - D_n \Sigma_{A_{mn}}^\infty P_m^\top)$.
1529
1530

1531 *Proof.* Convergence of Γ_{mn}^t : From Lemma 6.3, we have the recursion as follows,
1532

$$1533 \Gamma_{mn}^{t+1} = D_n^t \Gamma_{mn}^t + 2\gamma B_{mn}^t \Sigma_{h_m}^t \\ 1534$$

1535 Taking limits $t \rightarrow \infty$ and using Proposition 7.1 we obtain,
1536

$$1537 \Gamma_{mn}^\infty = D_n \Gamma_{mn}^\infty + 2\gamma B_{mn} \Sigma_{h_m}^\infty \\ 1538 (I - D_n) \Gamma_{mn}^\infty = 2\gamma B_{mn} \Sigma_{h_m}^\infty \\ 1539$$

1540 Since $\rho(D_n) < 1$, the matrix $(I - D_n)$ is invertible, giving:
1541

$$1542 \Gamma_{mn}^\infty = (I - D_n)^{-1} 2\gamma B_{mn} \Sigma_{h_m}^\infty \\ 1543$$

1544 Convergence of Ψ_{mn}^t : From Lemma 6.4, the recursion is given by,
1545

$$1546 \Psi_{mn}^{t+1} = D_n^t \Psi_{mn}^t H_m^{t\top} + D_n^t \Gamma_{mn}^t G_m^{t\top} - D_n^t \Sigma_{A_{mn}}^t P_m^{t\top} \\ 1547 + 2\gamma B_{mn}^t \Lambda_m^t H_m^{t\top} + 2\gamma B_{mn}^t \Sigma_{h_m}^t G_m^{t\top} - 2\gamma B_{mn}^t \Gamma_{mn}^{t\top} P_m^{t\top} \\ 1548$$

1549 At steady-state, using Proposition 7.1 we obtain,
1550

$$1551 \Psi_{mn}^\infty = D_n \Psi_{mn}^\infty H_m^\top + D_n \Gamma_{mn}^\infty G_m^\top - D_n \Sigma_{A_{mn}}^\infty P_m^\top \\ 1552 + 2\gamma B_{mn} \Lambda_m^\infty H_m^\top + 2\gamma B_{mn} \Sigma_{h_m}^\infty G_m^\top - 2\gamma B_{mn} \Gamma_{mn}^{\infty\top} P_m^\top \\ 1553$$

1554 This can be rewritten as a vectorized equation using $\text{Vec}(\cdot)$,
1555

$$1556 \text{Vec}(\Psi_{mn}^\infty) = (H_m \otimes D_n) \text{Vec}(\Psi_{mn}^\infty) + \text{Vec}(X) \\ 1557$$

1558 where X collects all remaining terms.

1559 Since $\rho(H_m \otimes D_n) = \rho(H_m) \rho(D_n) < 1$ by assumption, we have,
1560

$$1561 \text{Vec}(\Psi_{mn}^\infty) = (I - H_m \otimes D_n)^{-1} \text{Vec}(X) \\ 1562$$

1563 Substituting back X we obtain,
1564

$$1565 \Psi_{mn}^\infty = (I - H_m \otimes D_n)^{-1} \text{Vec}(D_n \Gamma_{mn}^\infty G_m^\top - D_n \Sigma_{A_{mn}}^\infty P_m^\top) \\ 1566$$

□

F.11 COROLLARY 7.3

• **Corollary** *The above assumptions lead to convergence of the uncertainty of the client states $\Sigma_{h_m}^t$ as follows: $\lim_{t \rightarrow \infty} \Sigma_{h_m}^t := \Sigma_{h_m}^\infty = \kappa_m \Sigma_{\theta_m}^\infty + \Omega_m^\infty (\mu_{y_m} \otimes I)^\top + (\mu_{y_m} \otimes I) \Omega_m^{\infty \top}$, where $\kappa_m = \text{tr}(\Sigma_{y_m}) + \|\mu_{y_m}\|^2$ and $\Omega_m^\infty = \Sigma_{\theta_m}^\infty \mu_{y_m}$.*

Proof. From Lemma 6.5, the client state variance evolves as follows,

$$\Sigma_{h_m}^t = \kappa_m^t \Sigma_{\theta_m}^t + \Omega_m^t (\mu_{y_m}^t \otimes I_{p_m})^\top + (\mu_{y_m}^t \otimes I_{p_m}) \Omega_m^{t \top}$$

where $\kappa_m^t = \text{tr}(\Sigma_{y_m}^t) + \|\mu_{y_m}^t\|^2$.

Under the stationarity Assumption (A4) and Proposition 7.1 we define the following,

$$\begin{aligned} \Sigma_{\theta_m}^\infty &:= \lim_{t \rightarrow \infty} \Sigma_{\theta_m}^t \\ \Omega_m^\infty &:= \lim_{t \rightarrow \infty} \Omega_m^t \\ \kappa_m &:= \lim_{t \rightarrow \infty} \kappa_m^t = \text{tr}(\Sigma_{y_m}) + \|\mu_{y_m}\|^2 \end{aligned}$$

From Proposition 6.1 and the steady-state analysis we have,

$$\Omega_m^\infty = \text{Cov}(v_m^\infty, y_m) = \Sigma_{\theta_m}^\infty \mu_{y_m}$$

since at steady-state, the parameter covariance dominates the data-model correlation.

Substituting these limits into the variance expression, we obtain,

$$\begin{aligned} \Sigma_{h_m}^\infty &= \kappa_m \Sigma_{\theta_m}^\infty + \Sigma_{\theta_m}^\infty \mu_{y_m} (\mu_{y_m} \otimes I_{p_m})^\top + (\mu_{y_m} \otimes I_{p_m}) \mu_{y_m}^\top \Sigma_{\theta_m}^\infty \\ &= \kappa_m \Sigma_{\theta_m}^\infty + \Sigma_{\theta_m}^\infty (\mu_{y_m} \mu_{y_m}^\top \otimes I_{p_m}) + (\mu_{y_m} \mu_{y_m}^\top \otimes I_{p_m}) \Sigma_{\theta_m}^\infty \end{aligned}$$

Simplifying using the Kronecker product properties, we have,

$$\lim_{t \rightarrow \infty} \Sigma_{h_m}^t = \kappa_m \Sigma_{\theta_m}^\infty + \Omega_m^\infty (\mu_{y_m} \otimes I_{p_m})^\top + (\mu_{y_m} \otimes I_{p_m}) \Omega_m^{\infty \top}$$

with $\Omega_m^\infty = \Sigma_{\theta_m}^\infty \mu_{y_m}$. □

F.12 THEOREM 7.4

• **Theorem (Convergence of Server Model’s Uncertainty)** Let $\rho(D_n) < 1$. Define the linear map $\mathcal{L}_n(X) = D_n X D_n^\top$ and the injection $Q_{mn}(\Sigma) = 4\gamma^2 B_{mn} (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) B_{mn}^\top$. Then, $\Sigma_{A_{mn}}^\infty := \lim_{t \rightarrow \infty} \Sigma_{A_{mn}}^t$ exists, is unique, and is given by, $\Sigma_{A_{mn}}^\infty = \sum_{k=0}^\infty \mathcal{L}_n^k(Q_{mn}(\Sigma_{\theta_m}^\infty))$

Proof. From Theorem 6.7, the server parameter covariance evolves as follows,

$$\begin{aligned} \Sigma_{A_{mn}}^{t+1} &= D_n^t \Sigma_{A_{mn}}^t D_n^{t \top} + 4\gamma^2 B_{mn}^t \Sigma_{h_m}^t B_{mn}^{t \top} \\ &\quad + 2\gamma (D_n^t \Gamma_{mn}^t B_{mn}^{t \top} + B_{mn}^t \Gamma_{mn}^{t \top} D_n^{t \top}) \end{aligned}$$

Taking limits $t \rightarrow \infty$ and using Proposition 7.1 we have,

$$\begin{aligned} \Sigma_{A_{mn}}^\infty &= D_n \Sigma_{A_{mn}}^\infty D_n^\top + 4\gamma^2 B_{mn} \Sigma_{h_m}^\infty B_{mn}^\top \\ &\quad + 2\gamma (D_n \Gamma_{mn}^\infty B_{mn}^\top + B_{mn} \Gamma_{mn}^{\infty \top} D_n^\top) \end{aligned}$$

From Corollary 7.3, we substitute $\Sigma_{h_m}^\infty$ as follows,

$$\Sigma_{h_m}^\infty = \kappa_m \Sigma_{\theta_m}^\infty + \Sigma_{\theta_m}^\infty M_m M_m^\top + M_m M_m^\top \Sigma_{\theta_m}^\infty$$

where $M_m = \mu_{y_m} \otimes I_{p_m}$ and $\kappa_m = \text{tr}(\Sigma_{y_m}) + \|\mu_{y_m}\|^2$.

Defining the linear operator $\mathcal{L}_n(X) = D_n X D_n^\top$ and the quadratic form we have,

$$Q_{mn}(\Sigma) = 4\gamma^2 B_{mn} (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) B_{mn}^\top$$

The steady-state equation thus becomes,

$$\Sigma_{A_{mn}}^\infty = \mathcal{L}_n(\Sigma_{A_{mn}}^\infty) + Q_{mn}(\Sigma_{\theta_m}^\infty)$$

Since $\rho(D_n) < 1$ (given as a condition), the operator $\mathcal{L}_n(\cdot)$ is a contraction, and the solution is given by the Neumann series:

$$\Sigma_{A_{mn}}^\infty = \sum_{k=0}^{\infty} \mathcal{L}_n^k(Q_{mn}(\Sigma_{\theta_m}^\infty))$$

With I being the identity operator, the above expression can be re-written as,

$$(I - \mathcal{L}_n)(\Sigma_{A_{mn}}^\infty) = Q_{mn}$$

This has a formal solution with operator inversion such that,

$$\Sigma_{A_{mn}}^\infty = (I - \mathcal{L}_n)^{-1} Q_{mn}$$

Since $\rho(D_n) < 1$, the Neumann series expansion is valid thus we can use,

$$(I - \mathcal{L}_n)^{-1} = \sum_{k=0}^{\infty} \mathcal{L}_n^k$$

Substituting this into the equation for $\Sigma_{A_{mn}}^\infty$ we obtain,

$$\Sigma_{A_{mn}}^\infty = \sum_{k=0}^{\infty} \mathcal{L}_n^k(Q_{mn})$$

The series converges because $\|\mathcal{L}_n^k(Q_{mn})\| \leq \rho(D_n)^{2k} \|Q_{mn}\| \rightarrow 0$ as $k \rightarrow \infty$. Uniqueness follows from the Banach fixed-point theorem, as \mathcal{L}_n is a contraction mapping on the space of positive semidefinite matrices with a matrix norm. \square

F.13 THEOREM 7.5

• Theorem (Convergence of Client Model’s Uncertainty) Let $\rho(H_m) < 1$. Write $\mathcal{M}_m(\Sigma) = H_m \Sigma H_m^\top$ and $R_m(\Sigma) = G_m (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) G_m^\top$. Then the steady-state $\Sigma_{\theta_m}^\infty := \lim_{t \rightarrow \infty} \Sigma_{\theta_m}^t$ is the unique solution to $\Sigma_{\theta_m}^\infty = \mathcal{M}_m(\Sigma_{\theta_m}^\infty) + R_m(\Sigma_{\theta_m}^\infty) + P_m \Sigma_{A_{mn}}^\infty P_m^\top$.

Proof. From Theorem 6.8, the client parameter covariance evolves as follows,

$$\Sigma_{\theta_m}^{t+1} = H_m^t \Sigma_{\theta_m}^t H_m^{t\top} + R_m^t(\Sigma_{\theta_m}^t) + P_m^t \Sigma_{A_{mn}}^t P_m^{t\top} + \text{cross terms}$$

where R_m^t collects terms quadratic in $\Sigma_{\theta_m}^t$.

Under Proposition 7.1’s convergence and Theorem 7.4’s steady-state for $\Sigma_{A_{mn}}^\infty$, we take limits as,

$$\Sigma_{\theta_m}^\infty = H_m \Sigma_{\theta_m}^\infty H_m^\top + R_m(\Sigma_{\theta_m}^\infty) + P_m \Sigma_{A_{mn}}^\infty P_m^\top$$

The quadratic term R_m derives from Corollary 7.3’s expression as,

$$R_m(\Sigma) = G_m (\kappa_m \Sigma + \Sigma M_m M_m^\top + M_m M_m^\top \Sigma) G_m^\top$$

with $G_m = 2\eta_1 (\mu_{y_m} \otimes (C_{mm} A_{mm})^\top)$ and $M_m = \mu_{y_m} \otimes I_{p_m}$.

Rewriting the fixed-point equation using the linear operator $\mathcal{M}_m(X) = H_m X H_m^\top$ we obtain,

$$\Sigma_{\theta_m}^\infty = \mathcal{M}_m(\Sigma_{\theta_m}^\infty) + R_m(\Sigma_{\theta_m}^\infty) + P_m \Sigma_{A_{mn}}^\infty P_m^\top$$

Since $\rho(H_m) < 1$ (given), \mathcal{M}_m is a contraction, guaranteeing a unique solution. \square

1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684

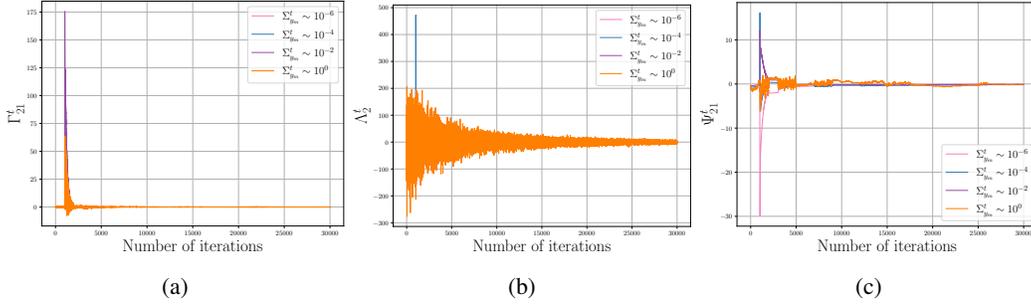


Figure 5: Uncertainty propagation in the cross-covariance terms during FedGC learning for different regimes of $\Sigma_{y_m}^t$ (a) Γ_{21}^t vs iterations, (b) Λ_2^t vs iterations, (c) Ψ_{21}^t vs iterations.

1688
1689
1690
1691
1692
1693
1694
1695
1696
1697

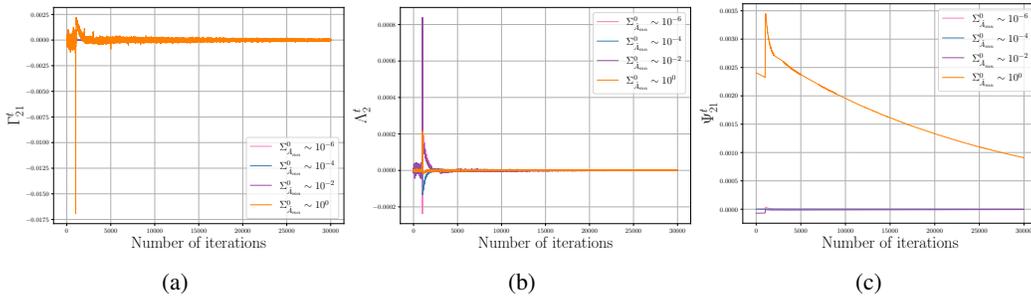


Figure 6: Uncertainty propagation in the cross-covariance terms during FedGC learning for different regimes of $\Sigma_{A_{mn}}^0$ (a) Γ_{21}^t vs iterations, (b) Λ_2^t vs iterations, (c) Ψ_{21}^t vs iterations.

1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711

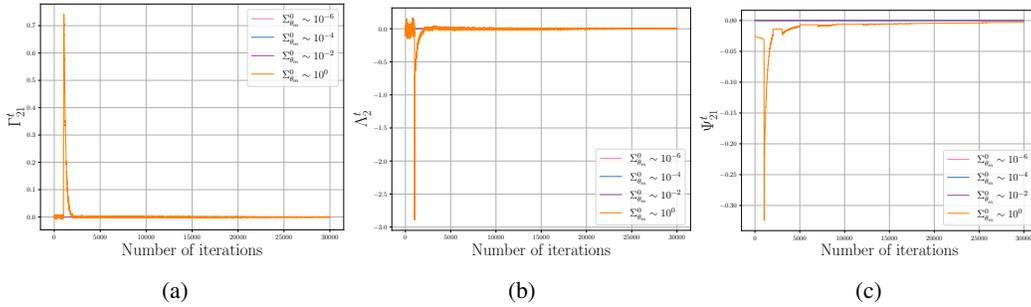


Figure 7: Uncertainty propagation in the cross-covariance terms during FedGC learning for different regimes of $\Sigma_{\theta_m}^t$ (a) Γ_{21}^t vs iterations, (b) Λ_2^t vs iterations, (c) Ψ_{21}^t vs iterations.

1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725

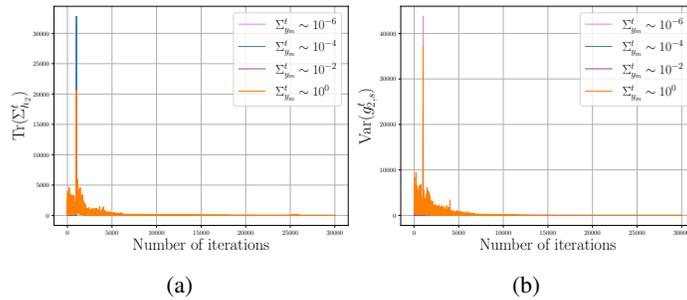


Figure 8: Uncertainty propagation in the communicated terms during FedGC learning for different regimes of $\Sigma_{y_m}^t$ (a) $\text{Var}(g_{m,s}^t)$ vs iterations, (b) $\Sigma_{h_m}^t$ vs iterations,

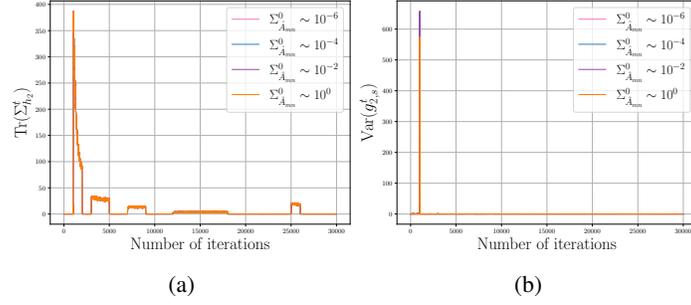


Figure 9: Uncertainty propagation in the communicated terms during FedGC learning for different regimes of $\Sigma_{\hat{A}_{mn}}^0$ (a) $\text{Var}(g_{m,s}^t)$ vs iterations, (b) $\Sigma_{h_m}^t$ vs iterations,

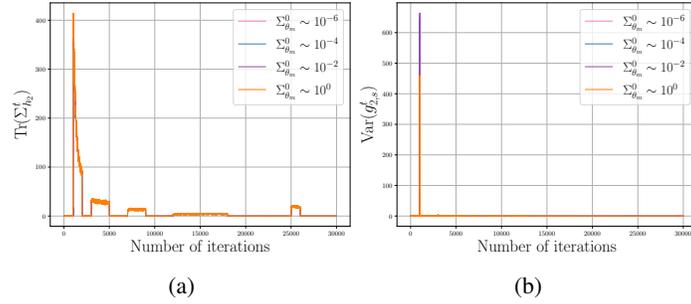


Figure 10: Uncertainty propagation in the communicated terms during FedGC learning for different regimes of $\Sigma_{\theta_m}^t$ (a) $\text{Var}(g_{m,s}^t)$ vs iterations, (b) $\Sigma_{h_m}^t$ vs iterations,

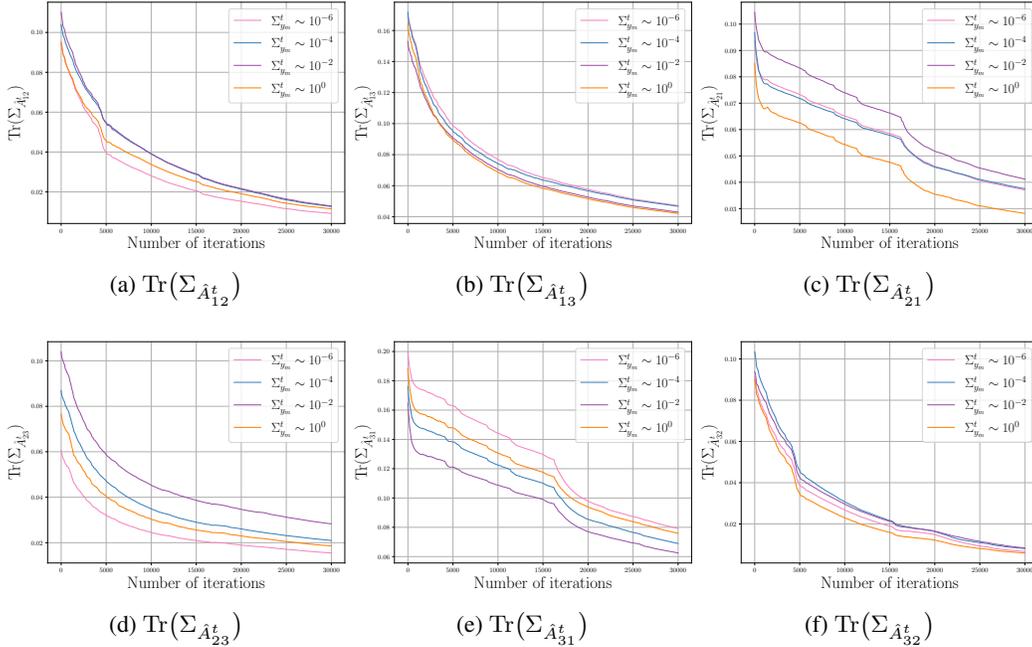


Figure 11: Trace of the covariance for each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{y_m}^t$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\hat{A}_{12}}^t)$, (b) $\text{Tr}(\Sigma_{\hat{A}_{13}}^t)$, (c) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (d) $\text{Tr}(\Sigma_{\hat{A}_{23}}^t)$, (e) $\text{Tr}(\Sigma_{\hat{A}_{31}}^t)$, (f) $\text{Tr}(\Sigma_{\hat{A}_{32}}^t)$ vs. iteration t .

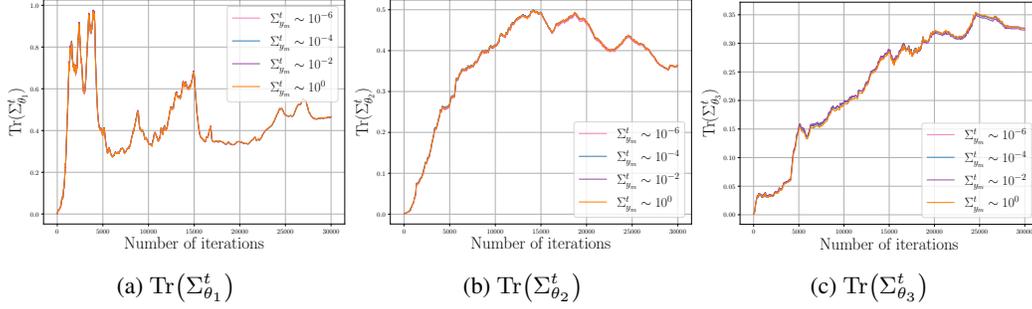


Figure 12: Trace of the covariance of the local model parameters θ for each of the three components during FedGC learning for different regimes of $\Sigma_{y_m}^t$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\theta_1}^t)$, (b) $\text{Tr}(\Sigma_{\theta_2}^t)$, (c) $\text{Tr}(\Sigma_{\theta_3}^t)$ vs. iteration t .

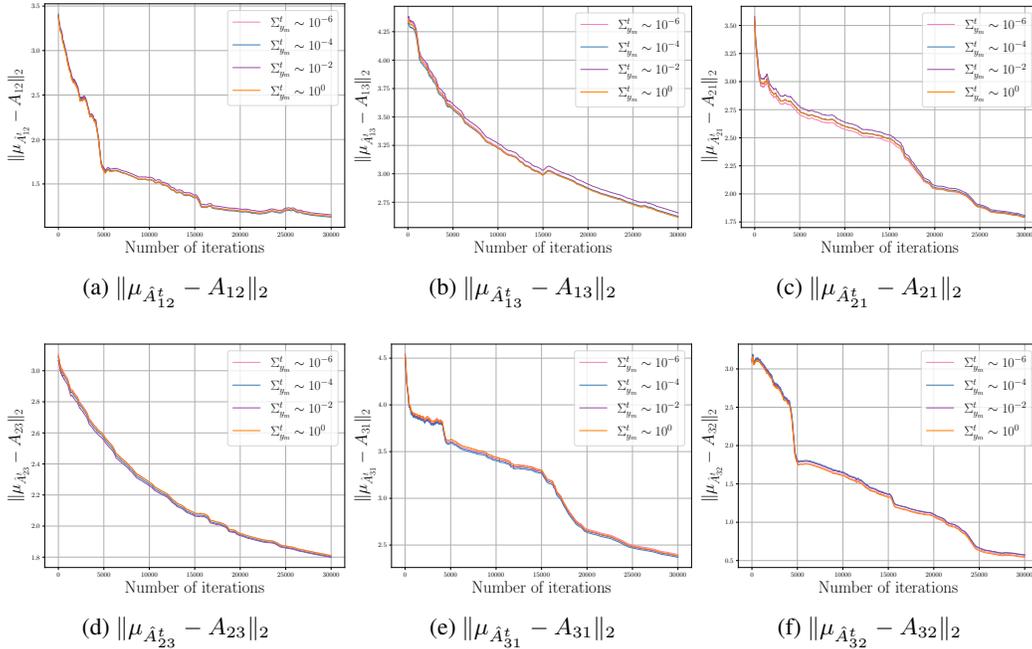


Figure 13: Average L_2 norm error of each off-diagonal block of the matrix A during FedGC learning for different regimes of $\Sigma_{y_m}^t$ for HAI dataset: (a) $\|\mu_{\hat{A}_{12}}^t - A_{12}\|_2$, (b) $\|\mu_{\hat{A}_{13}}^t - A_{13}\|_2$, (c) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$, (d) $\|\mu_{\hat{A}_{23}}^t - A_{23}\|_2$, (e) $\|\mu_{\hat{A}_{31}}^t - A_{31}\|_2$, (f) $\|\mu_{\hat{A}_{32}}^t - A_{32}\|_2$ vs. iteration t .

1836
 1837
 1838
 1839
 1840
 1841
 1842
 1843
 1844
 1845
 1846
 1847
 1848
 1849
 1850
 1851
 1852
 1853
 1854
 1855
 1856
 1857
 1858
 1859
 1860
 1861
 1862
 1863
 1864
 1865
 1866
 1867
 1868
 1869
 1870
 1871
 1872
 1873
 1874
 1875
 1876
 1877
 1878
 1879
 1880
 1881
 1882
 1883
 1884
 1885
 1886
 1887
 1888
 1889

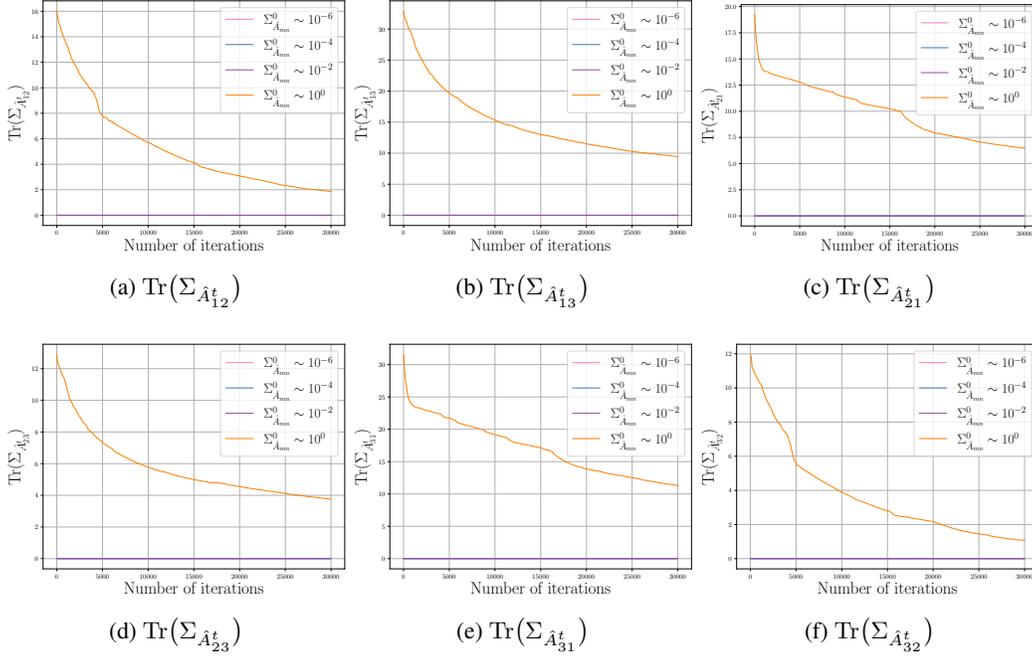


Figure 14: Trace of the covariance for each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{A_{mn}}^0$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\hat{A}_{12}}^t)$, (b) $\text{Tr}(\Sigma_{\hat{A}_{13}}^t)$, (c) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (d) $\text{Tr}(\Sigma_{\hat{A}_{23}}^t)$, (e) $\text{Tr}(\Sigma_{\hat{A}_{31}}^t)$, (f) $\text{Tr}(\Sigma_{\hat{A}_{32}}^t)$ vs. iteration t .

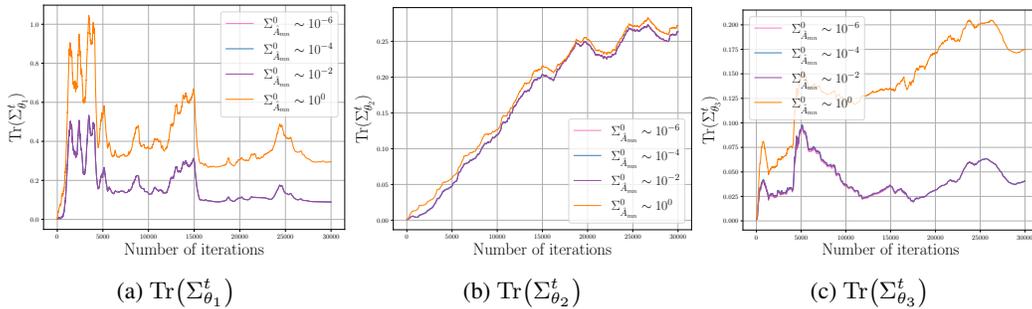


Figure 15: Trace of the covariance of the local model parameters θ for each of the three components during FedGC learning for different regimes of $\Sigma_{A_{mn}}^0$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\theta_1}^t)$, (b) $\text{Tr}(\Sigma_{\theta_2}^t)$, (c) $\text{Tr}(\Sigma_{\theta_3}^t)$ vs. iteration t .

1890
 1891
 1892
 1893
 1894
 1895
 1896
 1897
 1898
 1899
 1900
 1901
 1902
 1903
 1904
 1905
 1906
 1907
 1908
 1909
 1910
 1911
 1912
 1913
 1914
 1915
 1916
 1917
 1918
 1919
 1920
 1921
 1922
 1923
 1924
 1925
 1926
 1927
 1928
 1929
 1930
 1931
 1932
 1933
 1934
 1935
 1936
 1937
 1938
 1939
 1940
 1941
 1942
 1943

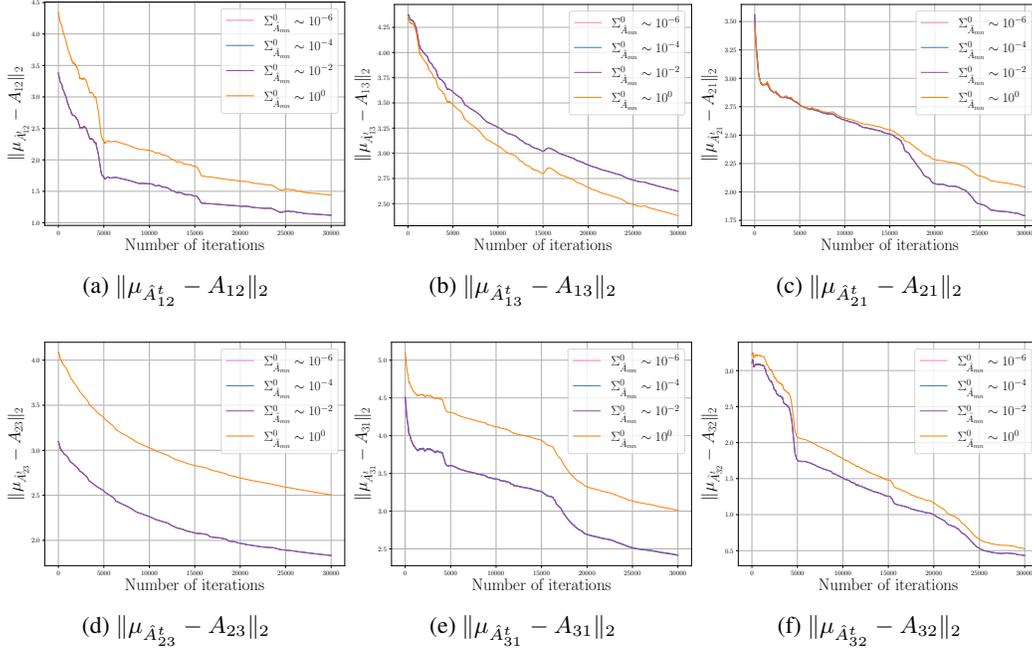


Figure 16: Average L_2 norm error of each off-diagonal block of the matrix A during FedGC learning for different regimes of $\Sigma_{\hat{A}_m}^0$ for HAI dataset: (a) $\|\mu_{\hat{A}_{12}}^t - A_{12}\|_2$, (b) $\|\mu_{\hat{A}_{13}}^t - A_{13}\|_2$, (c) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$, (d) $\|\mu_{\hat{A}_{23}}^t - A_{23}\|_2$, (e) $\|\mu_{\hat{A}_{31}}^t - A_{31}\|_2$, (f) $\|\mu_{\hat{A}_{32}}^t - A_{32}\|_2$ vs. iteration t .

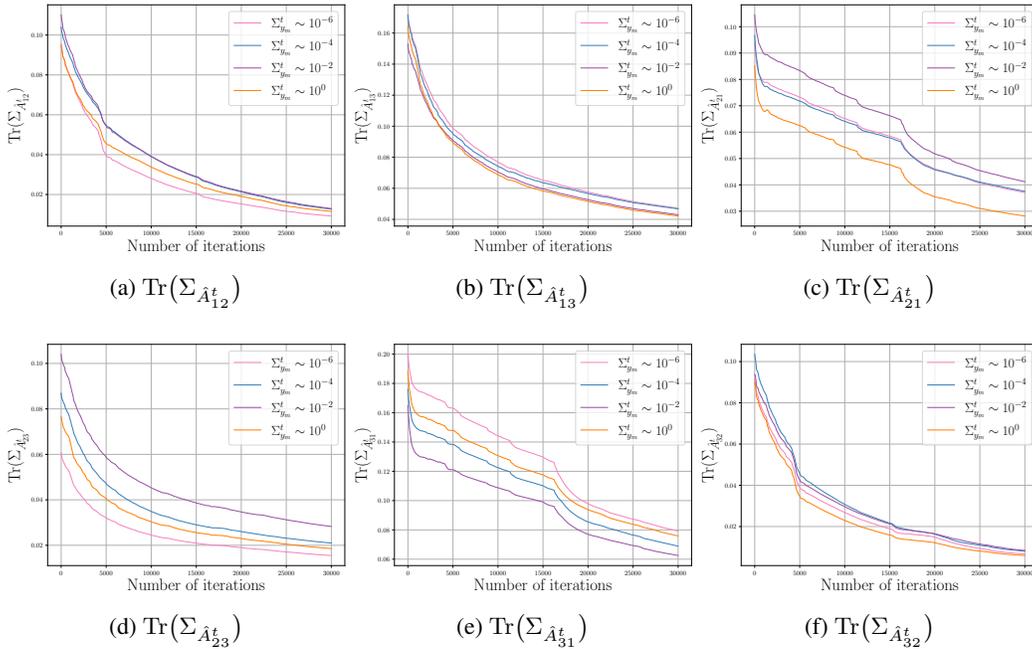


Figure 17: Trace of the covariance for each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\hat{A}_{12}}^t)$, (b) $\text{Tr}(\Sigma_{\hat{A}_{13}}^t)$, (c) $\text{Tr}(\Sigma_{\hat{A}_{21}}^t)$, (d) $\text{Tr}(\Sigma_{\hat{A}_{23}}^t)$, (e) $\text{Tr}(\Sigma_{\hat{A}_{31}}^t)$, (f) $\text{Tr}(\Sigma_{\hat{A}_{32}}^t)$ vs. iteration t .

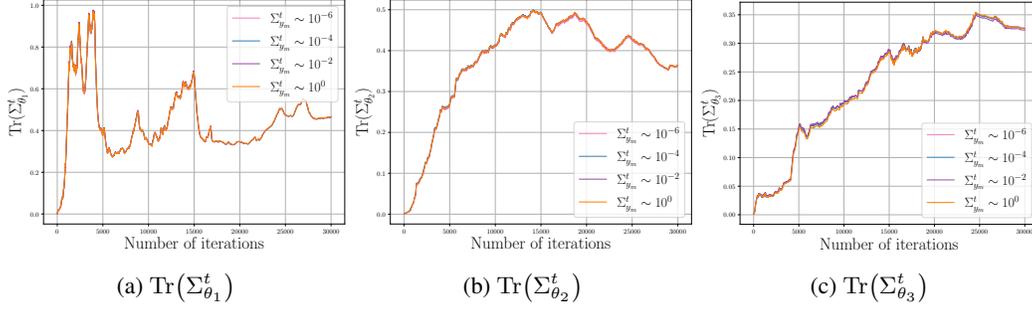


Figure 18: Trace of the covariance of the local model parameters θ for each of the three components during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ for HAI dataset: (a) $\text{Tr}(\Sigma_{\theta_1}^t)$, (b) $\text{Tr}(\Sigma_{\theta_2}^t)$, (c) $\text{Tr}(\Sigma_{\theta_3}^t)$ vs. iteration t .

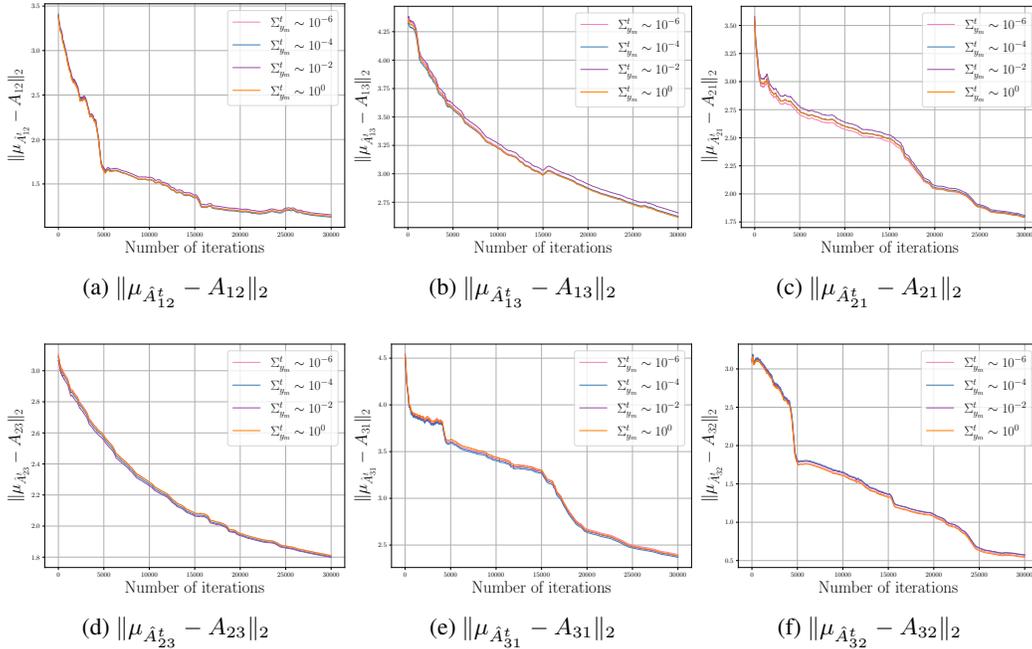
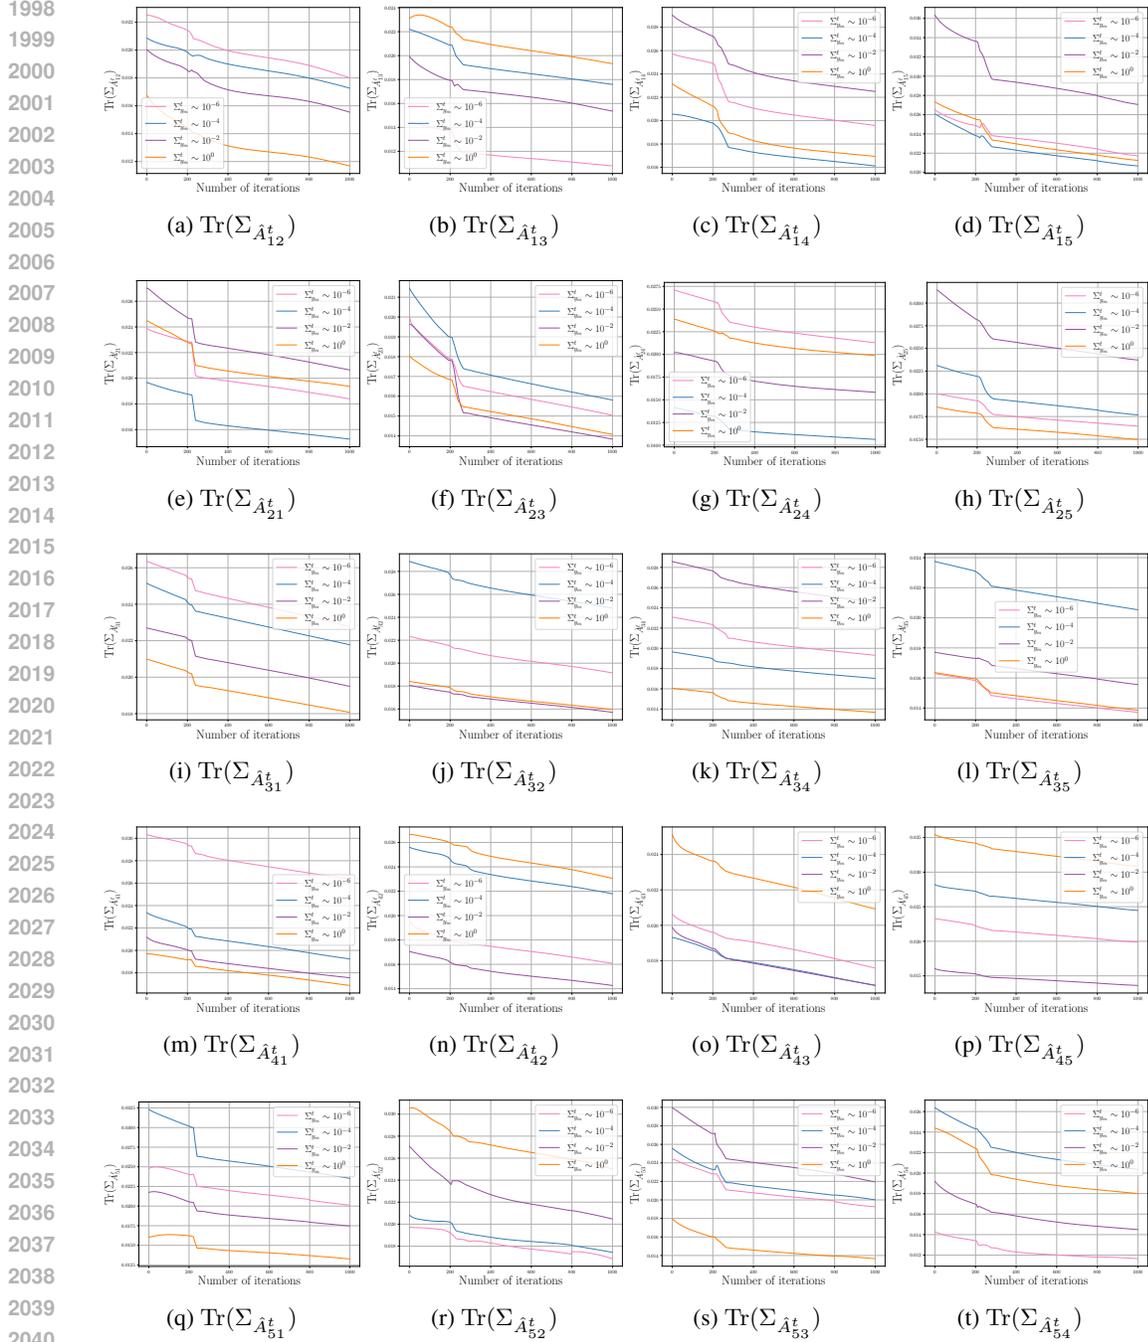
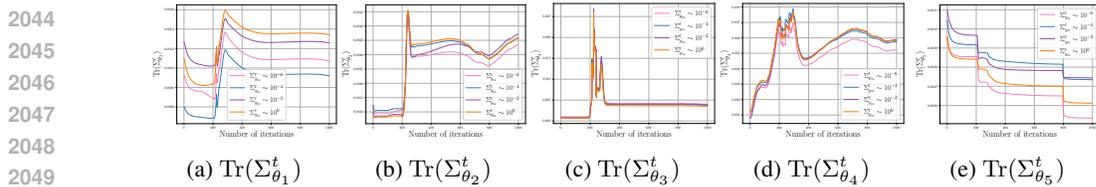


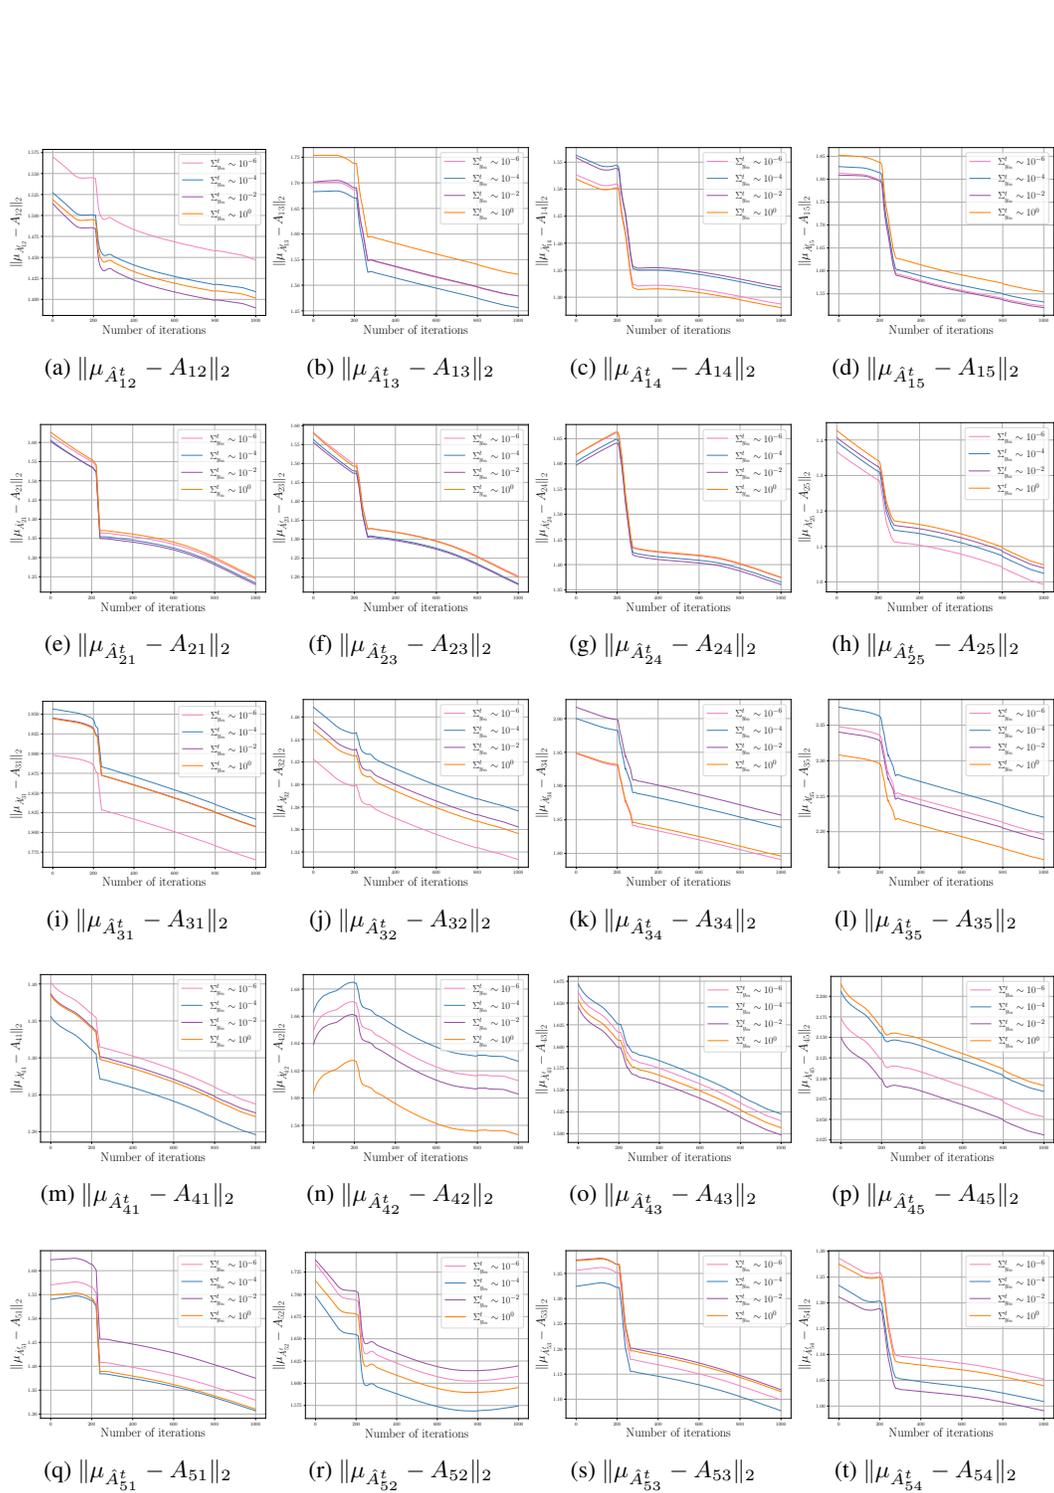
Figure 19: Average L_2 norm error of each off-diagonal block of the matrix A during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ for HAI dataset: (a) $\|\mu_{\hat{A}_{12}}^t - A_{12}\|_2$, (b) $\|\mu_{\hat{A}_{13}}^t - A_{13}\|_2$, (c) $\|\mu_{\hat{A}_{21}}^t - A_{21}\|_2$, (d) $\|\mu_{\hat{A}_{23}}^t - A_{23}\|_2$, (e) $\|\mu_{\hat{A}_{31}}^t - A_{31}\|_2$, (f) $\|\mu_{\hat{A}_{32}}^t - A_{32}\|_2$ vs. iteration t .



2041 **Figure 20:** Trace of the covariance for each off-diagonal block of the A matrix during FedGC learn-
2042 ing for different regimes of $\Sigma_{y_m}^t$ on SWaT dataset, plotted vs. iteration t .



2050 **Figure 21:** Trace of the covariance of the local model parameters θ_m for each of the five components
2051 during FedGC learning for different regimes of $\Sigma_{y_m}^t$ on SWaT dataset, plotted vs. iteration t .



2100 Figure 22: Average L_2 norm error of each off-diagonal block of the 5×5 matrix A during FedGC
2101 learning for different regimes of $\Sigma_{y_m}^t$ on SwaT dataset, plotted vs. iteration t .
2102
2103
2104
2105

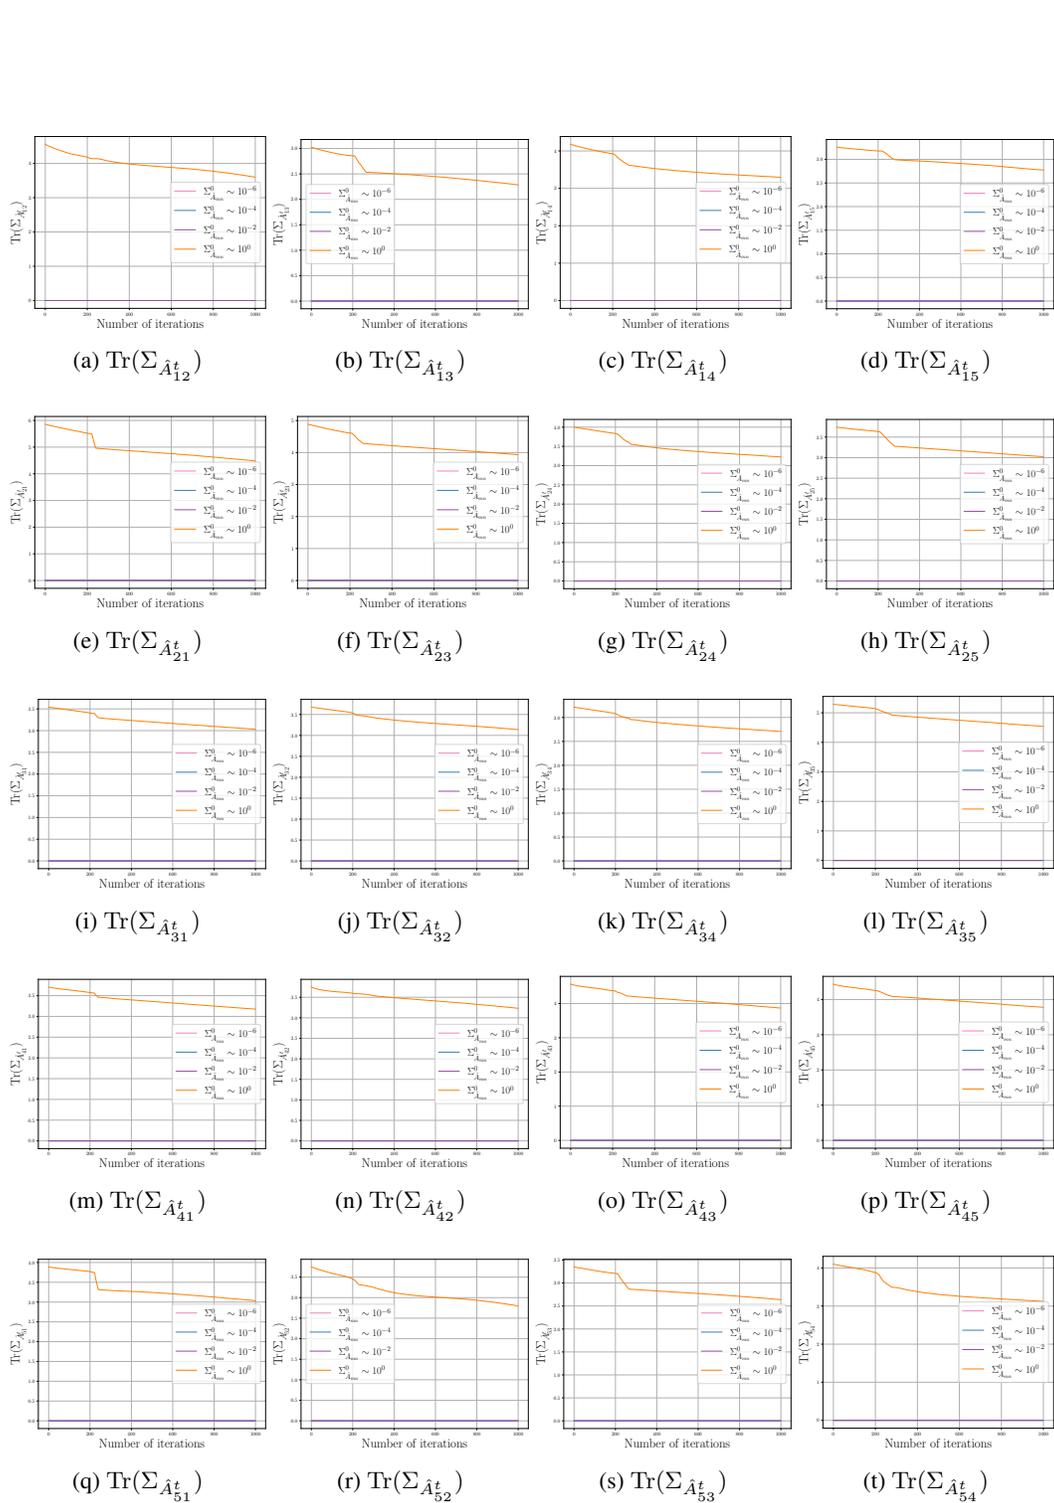


Figure 23: Trace of the covariance for each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{A_{mn}}^0$ on SwAT dataset, plotted vs. iteration t .

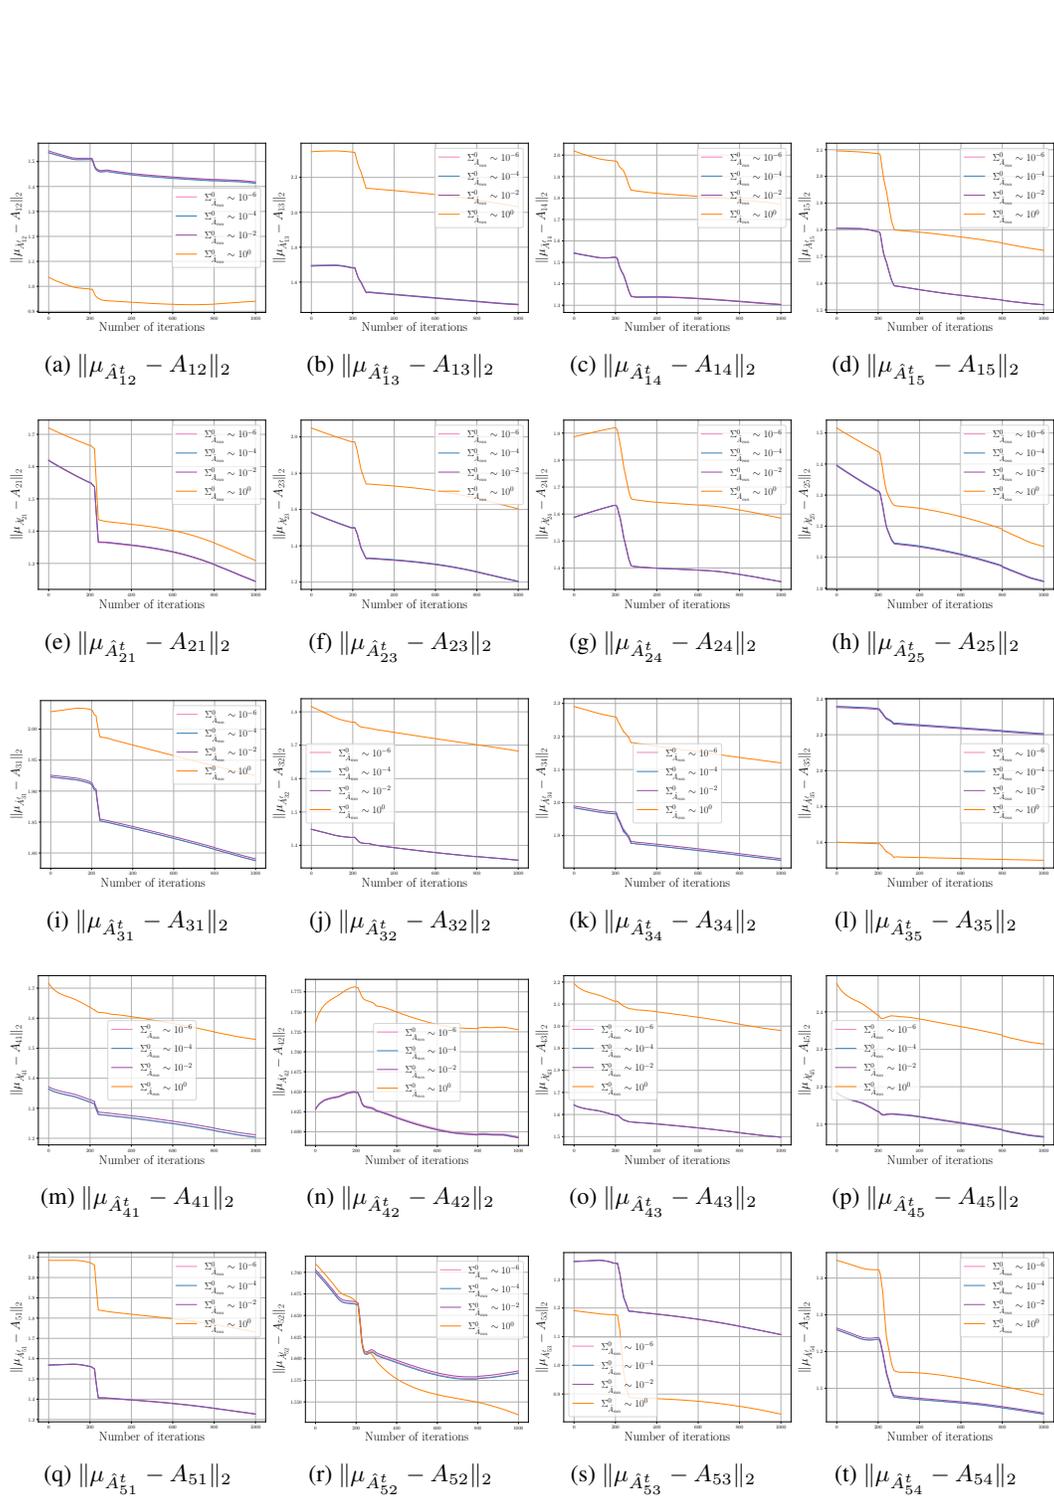


Figure 24: Average L_2 norm error of each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{\hat{A}_{mn}}^0$ on SWaT dataset, plotted vs. iteration t .

2214
 2215
 2216
 2217
 2218
 2219
 2220
 2221
 2222
 2223
 2224
 2225
 2226
 2227
 2228
 2229
 2230
 2231
 2232
 2233
 2234
 2235
 2236
 2237
 2238
 2239
 2240
 2241
 2242
 2243
 2244
 2245
 2246
 2247
 2248
 2249
 2250
 2251
 2252
 2253
 2254
 2255
 2256
 2257
 2258
 2259
 2260
 2261
 2262
 2263
 2264
 2265
 2266
 2267

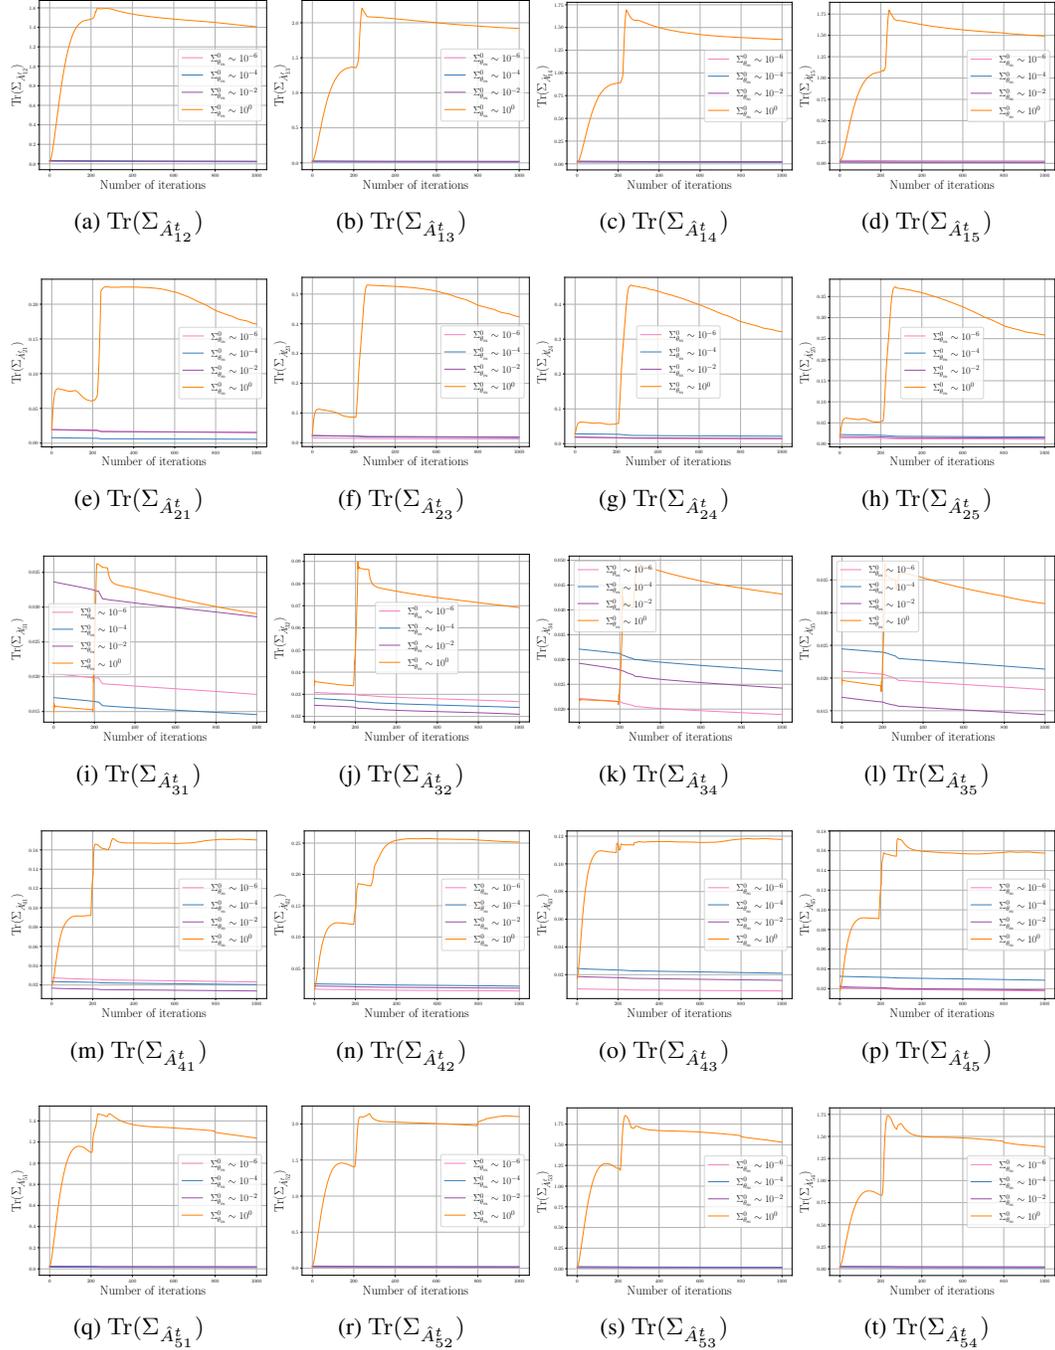


Figure 25: Trace of the covariance for each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ on SWaT dataset, plotted vs. iteration t .

2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321

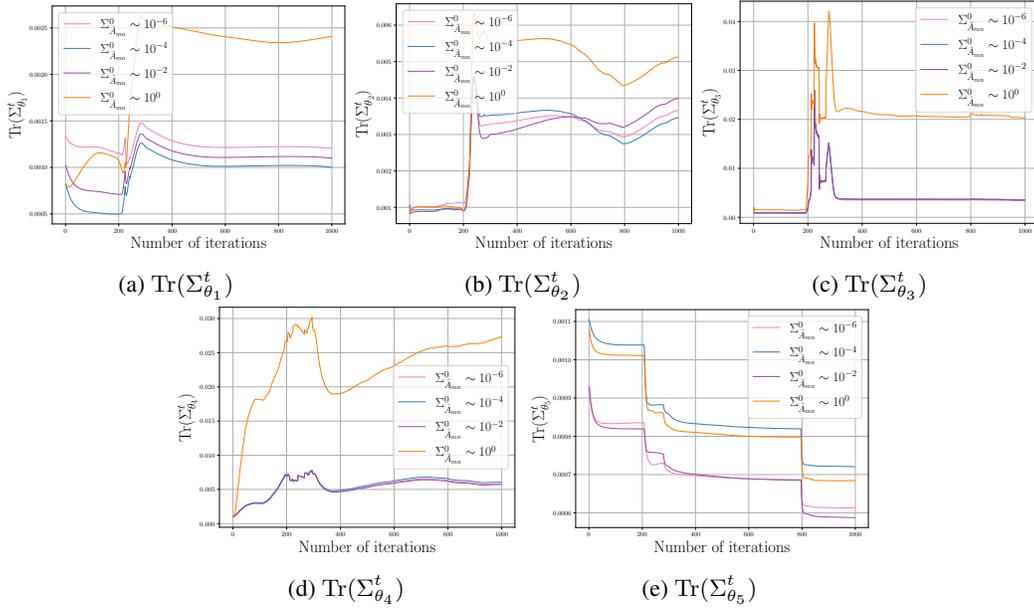


Figure 26: Trace of the covariance of the local model parameters θ for each of the five components during FedGC learning for different regimes of $\Sigma_{A_{mn}}^0$ on SWaT dataset, plotted vs. iteration t .

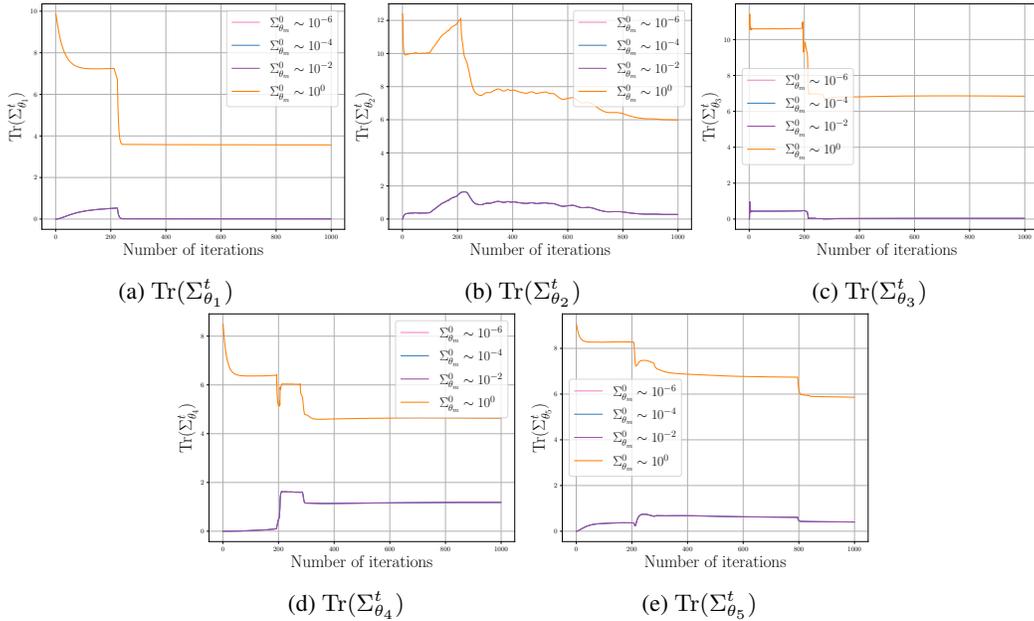


Figure 27: Trace of the covariance of the local model parameters θ for each of the five components during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ on SWaT dataset, plotted vs. iteration t .

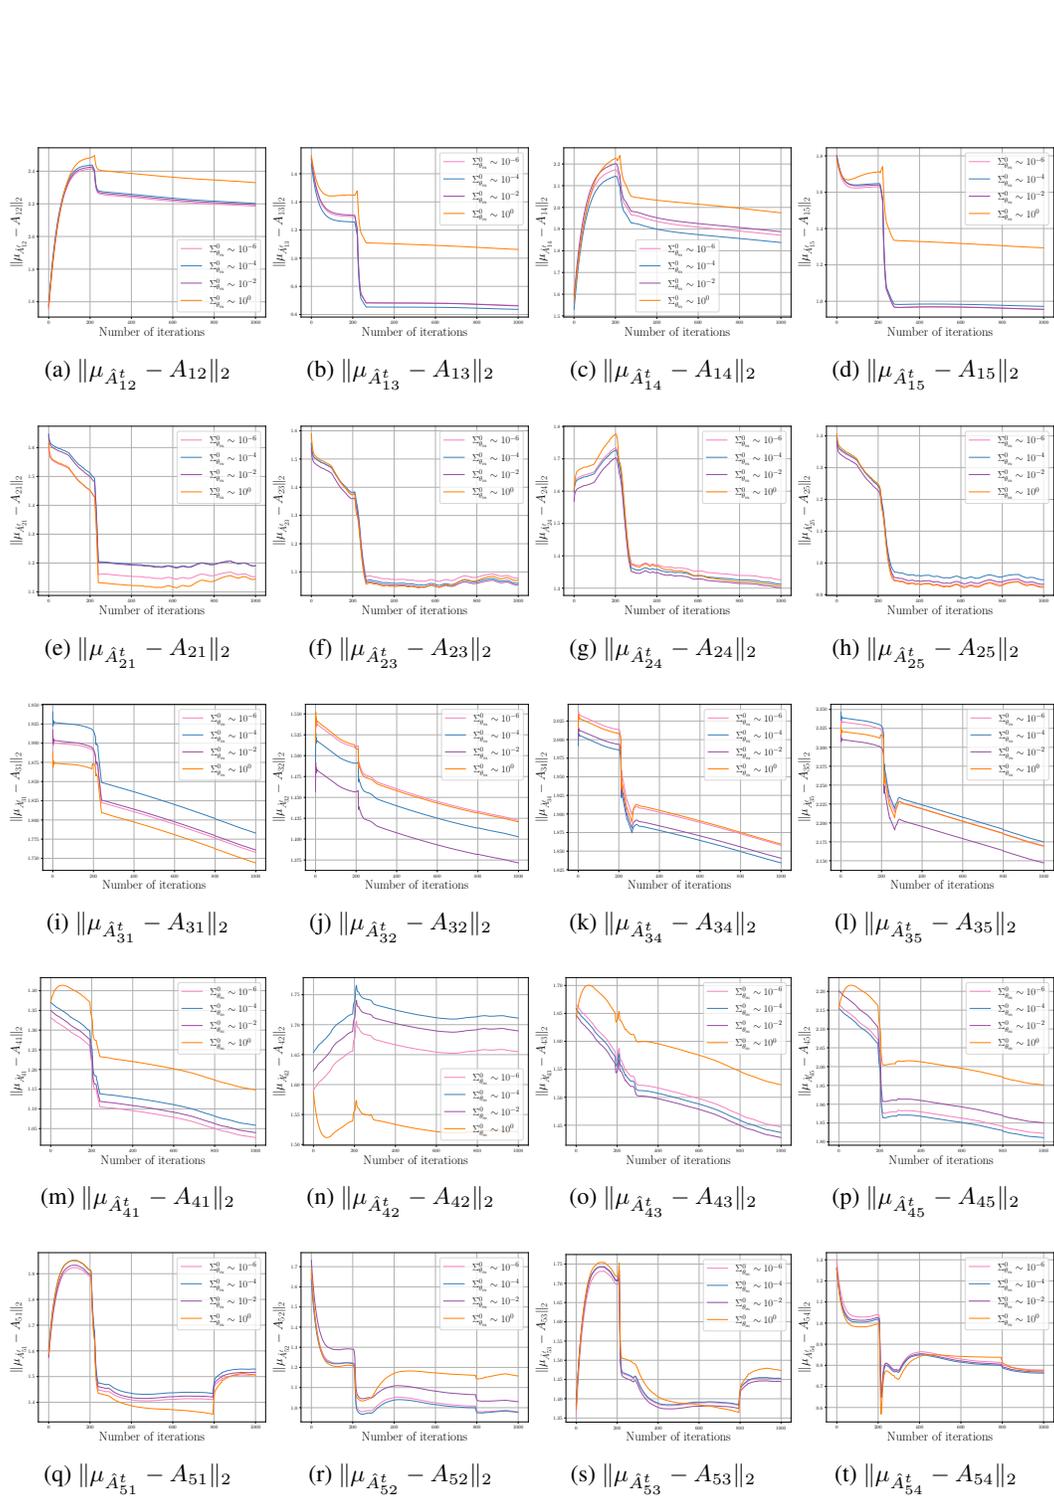


Figure 28: Average L_2 norm error of each off-diagonal block of the A matrix during FedGC learning for different regimes of $\Sigma_{\theta_m}^0$ on SWaT dataset, plotted vs. iteration t .