PRIVACY-AWARE DIRICHLET PROCESS MIXTURE MODELS: DIFFUSION-BASED TRAJECTORY PROTECTION AND DIFFERENTIALLY PRIVATE WEIGHTS

Anonymous authors
Paper under double-blind review

ABSTRACT

Mixture models are widely used to capture cluster structure in modern machine learning systems, particularly in time-varying applications such as smart grids and interacting large language models. However, releasing data from such systems can leak sensitive group information because the temporal trajectories of cluster centroids enable reconstruction of underlying group behavior. Static perturbation methods are insufficient, as repeated noise patterns can be filtered over time. We introduce a dynamic privacy framework based on a Dirichlet process mixture model (DDirPMM), in which centroids evolve through a diffusion process. The process injects temporally varying randomness to prevent trajectory reconstruction while preserving utility. In addition, our DDirPMM inherently induces static differential privacy for mixture weights, automatically protecting them without additional mechanisms. These two mechanisms address structurally distinct threat models, trajectory reconstruction and membership inference, respectively. Theoretical and empirical results demonstrate greater performance against filtering-based reconstruction and improved privacy-utility trade-offs compared with static trajectory protection methods.

1 INTRODUCTION

Mixture models provide a principled framework for capturing data structure through latent group centroids and their associated weights. Gaussian mixture models (GMMs) are a well-established approach for modeling complex, multimodal distributions and enabling probabilistic clustering (Viroli & McLachlan, 2019). They are widely used in smart grid applications for energy resource modeling and network equipment analysis, capturing cluster-wise time-varying patterns in electrical systems (Najafi & Livani, 2025; Li et al., 2025; Gemine et al., 2016). Beyond these applications, GMMs have recently been used to capture structure in large language models (LLMs) (Wang et al., 2025; Xing et al., 2025; Muttakhiroh & Fevens, 2025; Xu et al., 2025; Su et al., 2026). These developments naturally lead to dynamic mixture models with time-varying component parameters.

These dynamic settings introduce new privacy risks. GMM parameters directly encode group-level information over time: weights describe how populations are distributed across groups, while centroids capture latent group characteristics. When these parameters are released sequentially, their temporal evolution can reveal how groups are organized and how they change over time, enabling reconstruction of sensitive group-level patterns (Xiong et al., 2011; Chicco, 2012; Seo et al., 2025; Dai et al., 2025; Adam et al., 2021; Zhang et al., 2025).

In particular, centroids are especially vulnerable, as they can reveal sensitive aspects of latent group behavior through their trajectories over time (Sharma et al., 2025). A natural idea is to add independent noise at each release; however, such perturbations can be filtered out across repeated releases, enabling recovery of the underlying centroids.

Moreover, adding noise through a static structure across time imposes a uniform pattern of perturbation. This is misaligned with dynamic mixture models, where the informativeness of parameters varies over time. For example, periods of rapid centroid movement correspond to *sensitive time intervals*, as they reflect shifts in the underlying population structure, whereas stable periods convey

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

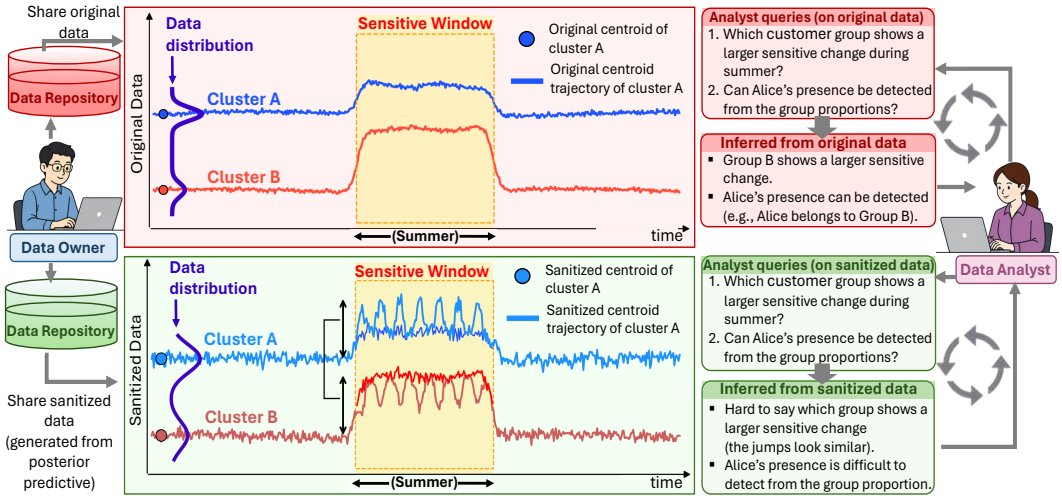

Figure 1: Overview of the privacy framework with a data-owner target mean. Example: protecting summer electricity profiles prevents unequal promotions from AC companies across customer groups.

less information. Applying the same perturbation across all time points therefore fails to provide stronger protection during sensitive times and introduces unnecessary utility loss elsewhere.

Our contribution. We introduce a time-varying privacy framework through a Bayesian nonparametric (BNP) mixture model, providing a privacy-preserving alternative to GMMs in dynamic settings. Specifically, we propose a dynamic Dirichlet process mixture model (DDirPMM) with a diffusion process (Oksendal, 2013) as a prior for the centroids. This allows their trajectories to evolve over time with structured stochasticity. This construction provides control over both the shape and magnitude of the temporal trajectories, protecting them from reconstruction while preserving statistical utility. This diffusion-based protection differs from well-known privacy notions such as differential privacy (DP), which protects against the addition or removal of an individual. We have not yet established whether this mechanism also provides neighboring-dataset indistinguishability. The DDirPMM is based on a finite representation of the Dirichlet process (Ishwaran & Zarepour, 2002), which additionally induces a Dirichlet-based DP mechanism for the mixture weights. The resulting model is defined through its posterior predictive distribution, combining time-evolving protected centroids with differentially private weights to generate synthetic data for release.

The remainder of the paper is organized as follows. Sections 2 and 3 review related work and background. Section 4 develops the framework and its privacy mechanisms. Section 5 evaluates it on a toy example and a smart meter dataset, showing improved protection against reconstruction and privacy-utility trade-offs. Section 6 concludes with extensions, including interacting LLM agents. A notational glossary, proofs, additional experiments, and algorithms are provided in the appendix.

2 Related Work

Most privacy-preserving clustering methods are based on DP (Wu et al., 2016; Liu et al., 2025; Tong et al., 2017; Guan et al., 2020; Ravi et al., 2022), with relatively few non-DP approaches (Li et al., 2022; 2020; Wu, 2025), none of which address time-varying privacy concerns. DP methods are designed to protect individual-level information by limiting the influence of any single observation. However, they are not well-suited for protecting the collective behavior of groups over time, as reflected in centroid trajectories (Xiong et al., 2011; Chicco, 2012; Seo et al., 2025; Dai et al., 2025).

DP mechanism. Wu et al. (2016) perturbs GMM parameters with independent noise and evaluates utility by mean squared error (MSE). Liu et al. (2025) follows this approach but uses Kullback–Leibler divergence to improve distributional fidelity. Tong et al. (2017) perturbs queried means and covariances, allowing noise to propagate through the GMM. However, this may distort parameters, e.g., produce non-positive definite covariances. Guan et al. (2020) estimates and then perturbs mixture parameters in a BNP model, where posterior uncertainty may provide additional privacy, though its benefit beyond the injected noise remains unclear. Beraha et al. (2025) takes a different direction, applying BNP mixture models to already DP-privatized data for density recovery. This highlights

that privatizing the data does not directly protect the governing mixture parameters. As a GMM alternative, Ravi et al. (2022) proposes a K-means-based method that adds position-dependent noise to centroids and cluster labels. The method is argued to better resist filtering-based recovery of true centroids, although this claim is not thoroughly evaluated.

Non-DP mechanism. Beyond DP, which protects the presence of individual records, Li et al. (2022) proposes a method to protect information shared when several participants jointly fit a GMM. Wu (2025) instead uses Pufferfish privacy to protect specified information about individuals in GMMs.

3 Background and Preliminaries

This section reviews key background on Dirichlet processes (DirP), their mixture models, differential privacy (DP), and mean-reverting Ornstein–Uhlenbeck processes (MR-OUP).

DirP and Mixture Representation. Let G be a random probability measure on a measurable space $(\mathfrak{X}, \mathcal{A})$. The DirP, denoted by $G \sim \mathrm{DP}(\alpha, \mathring{G})$, defines a prior over such random measures (Ferguson, 1973). For any measurable partition $A_1, \ldots, A_L$, it satisfies $(G(A_1), \ldots, G(A_L)) \sim \mathrm{Dir}(\alpha\mathring{G}(A_1), \ldots, \alpha\mathring{G}(A_L))$, where the base measure $\mathring{G}$ reflects prior beliefs and $\alpha > 0$ regulates their influence. An explicit representation of the DirP is given by the stick-breaking construction (Sethuraman, 1994), $G(\cdot) = \sum_{i=1}^{\infty} \omega_i \, \mathsf{D}_{\boldsymbol{\theta}_i}(\cdot)$, where $\mathsf{D}_{\boldsymbol{\theta}_i}(\cdot)$ is the Dirac delta measure, showing that the DirP is almost surely discrete. Here, $\{\omega_i\}_{i=1}^{\infty}$ are weights obtained by sequentially breaking a unit-length stick into infinitely many parts, governed by α, and $\{\boldsymbol{\theta}_i\}_{i=1}^{\infty}$ are i.i.d. draws from $\mathring{G}$.

Let $\{\mathbf{X}_i\}_{i=1}^{n} \sim P$ be i.i.d. random variables. A DirP mixture model (DirPMM) places a DirP prior on the mixing distribution with a parametric kernel $\kappa(\cdot \mid \boldsymbol{\theta})$, and is defined hierarchically as

$$\mathbf{X}_i \mid \boldsymbol{\theta}_i \sim \kappa(\cdot \mid \boldsymbol{\theta}_i), \quad \boldsymbol{\theta}_i \mid G \sim G, \quad G \sim \mathrm{DP}(\alpha, \mathring{G}).$$

To model the data distribution, the DirPMM represents the density as $p(\mathbf{x}|G) = \int \kappa(\mathbf{x} \mid \boldsymbol{\theta}) \, G(\boldsymbol{d\theta})$. Under the stick-breaking representation (Blei & Jordan, 2006), this yields the infinite mixture form $p(\mathbf{x}|G) = \sum_{i=1}^{\infty} \omega_i \, \kappa(\mathbf{x} \mid \boldsymbol{\theta}_i)$. The primary inferential goal is to estimate the data density through the posterior predictive (PP) distribution, $p(\mathbf{x}_{n+1} \mid \mathbf{x}_{1:n}) = \int p(\mathbf{x}_{n+1}|G) \, p(G \mid \mathbf{x}_{1:n}) \, dG$.

(ϵ, δ)-DP Mechanism. DP ensures that the inclusion or exclusion of a single data point does not significantly alter the output distribution of a randomized algorithm (Dwork & Roth, 2014). A mechanism $\mathcal{M} : \mathfrak{X} \to \mathbb{R}^{d'}$ satisfies (ϵ, δ)-DP if for any neighboring datasets $\mathbf{X}, \mathbf{X}' \in \mathbb{R}^d$ and measurable set $E \subseteq \mathbb{R}^{d'}$, $\Pr[\mathcal{M}(\mathbf{X}; \epsilon, \delta) \in E] \leq e^{\epsilon} \Pr[\mathcal{M}(\mathbf{X}'; \epsilon, \delta) \in E] + \delta$, where ϵ controls the privacy-utility trade-off, and δ allows a small failure probability. If $\mathcal{M}$ satisfies $(\epsilon, 0)$-DP, it is referred to as ϵ-DP.

MR-OUP with time periodic coefficients. An MR-OUP is a continuous-time stochastic process in the class of diffusion processes (Oksendal, 2013). It is commonly used to model dynamics that fluctuate around a long-term equilibrium. Let $\{\boldsymbol{\mu}(s) : s \in \mathcal{T}\}$ denote such a process, defined by the Itô stochastic differential equation (SDE)

$$d\boldsymbol{\mu}(s) = \beta\big(\mathbf{m} - \boldsymbol{\mu}(s)\big) \, ds + \sigma(s) \, d\mathbf{W}(s), \tag{1}$$

where the coefficient $\beta > 0$ is the mean-reversion rate, $\mathbf{m} \in \mathbb{R}^d$ is the long-term mean, and $\mathbf{W}(s)$ is a standard d-dimensional Wiener process. The diffusion coefficient $\sigma(s)$ governs the structure of stochastic fluctuations and is assumed to be progressively measurable with $\int_0^s |\sigma(u)|^2 du < \infty$ almost surely for all finite s. The classical MR-OUP assumes a constant diffusion coefficient, while equation 1 extends it to a time-dependent, including periodic, coefficient (Da Prato & Lunardi, 2007).

4 Methodology: (α, β)-Privacy-Aware Framework for Clustering

To protect centroid trajectories in dynamic settings and prevent membership leakage through latent group weights, we first develop a DDirPMM for dynamic clustering.

4.1 DYNAMIC DIRICHLET PROCESS MIXTURE MODEL (DDIRPMM)

Let $\{\mathbf{X}_i(s) : s \in \mathcal{T}\}_{i=1}^n$ be i.i.d. stochastic processes with corresponding marginal distributions P_s. To develop a privacy-aware BNP counterpart of the GMM, we consider the hierarchical model

$$\mathbf{X}_i(s) \mid \boldsymbol{\mu}_i(s), \Sigma_i \sim \mathrm{N}(\boldsymbol{\mu}_i(s), \Sigma_i), \quad (\boldsymbol{\mu}_i(s), \Sigma_i) \mid G \sim G, \quad G \sim \mathrm{DP}(\alpha, \mathring{G}), \tag{2}$$

where the centroids $\{\boldsymbol{\mu}_i(s) : s \in \mathcal{T}\}_{i=1}^n$ are modeled as stochastic processes evolving over time, while Σ_i is kept fixed. The base measure is decomposed as

$$\mathring{G}(\boldsymbol{\mu}_i(s), \Sigma_i) = \mathring{G}'(\boldsymbol{\mu}_i(s) \mid \Sigma_i) \times \mathring{G}''(\Sigma_i), \tag{3}$$

where $\mathring{G}'$ and $\mathring{G}''$ are typically Gaussian and inverse-Wishart (IW) distributions, respectively. Hereafter, $\mathring{g}$, $\mathring{g}'$, and $\mathring{g}''$ will be used to denote the corresponding densities.

However, a key challenge is how to define $\mathring{G}'$ so that it incorporates structured stochastic perturbations while preserving a Gaussian form, thereby retaining the computational framework commonly used in Gaussian DirPMMs. To address this, we model the centroid trajectories through the MR-OUP defined by the Itô SDE equation 1. More specifically, the MR-OUP in equation 1 is applied component-wise with corresponding parameters $\mathbf{m}_i, \beta_i, \sigma_i(s)$. This introduces structured stochastic perturbations on centroid trajectories while preserving temporal dependence. Here, $\mathbf{m}_i$ is selected by the data owner to determine the desired long-term target of the perturbed trajectory. The following lemma shows that equation 1 admits a closed-form Gaussian solution with explicit parameters, allowing the induced Gaussian distribution to be incorporated into $\mathring{G}'$.

Lemma 1 *Let $\{\boldsymbol{\mu}_i(s) : s \in \mathcal{T}\}$ be an MR-OUP defined by equation 1. If $\boldsymbol{\mu}_i(s_0) \sim \mathrm{N}(\mathbf{m}_i(s_0), \Sigma_i)$ independently of $\{\mathbf{W}(s) : s \in \mathcal{T}\}$, then $\boldsymbol{\mu}_i(s) \mid \Sigma_i \sim \mathrm{N}(\mathbf{m}_i(s), \Sigma_i(s))$ for any $s > s_0$, where*

$$\mathbf{m}_i(s) = \mathbf{m}_i + e^{-(s-s_0)\beta_i}(\mathbf{m}_i(s_0) - \mathbf{m}_i), \quad \Sigma_i(s) = c(s, \beta_i)\Sigma_i + q(s, \beta_i)\mathrm{I}_d, \tag{4}$$

with $c(s, \beta_i) = e^{-2(s-s_0)\beta_i}$ and $q(s, \beta_i) = \int_{s_0}^s e^{-2(s-u)\beta_i}\sigma_i^2(u)\, du$.

The following corollary characterizes the stationary behavior of centroid trajectories.

Corollary 1 (Stationary State) *Under the assumptions of Lemma 1, suppose that $\sigma_i(s) \to \sigma_i > 0$ as $s \to \infty$. Then, $\boldsymbol{\mu}_i(s) \mid \Sigma_i \xrightarrow{d} \mathrm{N}\Big(\mathbf{m}_i, \frac{\sigma_i^2}{2\beta_i}\mathrm{I}_d\Big)$.*

Following Lemma 1, we take $\mathring{G}'$ in equation 3 as $\mathrm{N}(\mathbf{m}_i(s), \Sigma_i(s))$, together with $\mathring{G}''$ as $\mathrm{IW}(\Lambda_0, \nu_0)$ for fixed Λ_0, ν_0. Applying the resulting $\mathring{G}$ in equation 2, and using the truncated representation proposed by Ishwaran & Zarepour (2002) to approximate the DirP, yields the finite mixture model with N components, representing latent groups,

$$p_N(\mathbf{x}(s) \mid G) = \sum_{j=1}^N \omega_j \underbrace{\mathrm{N}(\mathbf{x}(s); \boldsymbol{\mu}_j(s), \Sigma_j)}_{j\text{-th component}}, \quad (\omega_1, \ldots, \omega_N) \sim \mathrm{Dir}(\alpha/N, \ldots, \alpha/N)\,. \tag{5}$$

The equation 5 is a truncated single-weights dependent DirPMM (Quintana et al., 2022), where weights are shared across time while components evolve stochastically. See Remark 3 in Appendix A.2 for an adaptive rule to select N. The following lemma derives its PP distribution, which forms the basis for generating new observations from the fitted model.

Lemma 2 *Under equation 5, let $Z_i(s) \in [\![N]\!]$ denote the latent allocation variable assigning $\mathbf{X}_i(s)$, $i \in [\![n]\!]$, to component j, with $\Pr(Z_i(s) = j \mid \omega_{1:N}) = \omega_j$. Define $I_j(s) = \{i : Z_i(s) = j\}$, $n_j(s) = \mathrm{card}(I_j(s))$, and $\mathbf{X}_{I_j}(s) = \{\mathbf{X}_i(s) : i \in I_j(s)\}$. Then, upon observing $\mathbf{x}_{1:n}(s)$ and $z_{1:n}(s)$, the PP is*

$$p_N\left(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)\right) = \sum_{j=1}^N \mathbb{E}(\omega_j \mid z_{1:n}(s))\, \dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)), \tag{6}$$

where $\dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)) = \iint \mathrm{N}(\mathbf{x}_{n+1}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j)\, \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))\, d\boldsymbol{\mu}_j(s)\, d\Sigma_j$.

The asymptotic behavior of equation 6 as $N \to \infty$ is established in Remark 2 (Appendix). The next proposition shows that $\mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))$ in $\dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s))$ is not always closed form, motivating variational inference (Wang & Blei, 2013) for its calculation (see Appendix A.5).

Proposition 1 *Under the specification of $\mathring{G}'$ as $\mathrm{N}(\mathbf{m}_j(s), \Sigma_j(s))$ and $\mathring{G}''$ as $\mathrm{IW}(\Lambda_0, \nu_0)$ in equation 3, the posterior density admits the decomposition*

$$\mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)) = \mathring{g}'(\boldsymbol{\mu}_j(s) \mid \Sigma_j, \mathbf{x}_{I_j}(s))\, \mathring{g}''(\Sigma_j \mid \mathbf{x}_{I_j}(s)), \tag{7}$$

where $\mathring{g}'(\boldsymbol{\mu}_j(s) \mid \Sigma_j, \mathbf{x}_{I_j}(s))$ is a Gaussian density with

$$\mathbf{m}_{j,n_j}(s) = \Sigma_{j,n_j}(s)\Big(\Sigma_j^{-1}(s)\mathbf{m}_j(s) + n_j(s)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s)\Big),\ \Sigma_{j,n_j}(s) = \Big(\Sigma_j^{-1}(s) + n_j(s)\Sigma_j^{-1}\Big)^{-1}, \tag{8}$$

for all $s \geq s_0$. Moreover, $\mathring{g}''(\Sigma_j \mid \mathbf{x}_{I_j}(s))$ belongs to the IW family only at $s_0 = 0$. Consequently, the chosen prior is conjugate and yields a closed-form posterior in equation 7 only at $s_0 = 0$.

4.2 $\mathcal{U}_s(\beta)$-DIFFUSION PRIVACY FOR CENTROID TRAJECTORIES

We now investigate the role of the proposed diffusion scheme in perturbing centroid trajectories. The following theorem shows how diffusion draws the posterior mean toward the prior mean and away from the empirical cluster mean, while bounding the associated posterior uncertainty.

Theorem 1 *Under equation 4 and equation 8, let $\Sigma_j = \varsigma_j^2 \mathrm{I}_d$. Then, given the observed cluster data $\mathbf{x}_{I_j}(s) = \{\mathbf{x}_i(s) : i \in I_j(s)\}$ at any fixed time $s \in \mathcal{T}$, the following hold:*

(i) The posterior means satisfy

$$\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \leq \mathcal{U}_s(\beta_j)\|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2, \tag{9}$$

$$\|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \leq (1 - \mathcal{U}_s(\beta_j))\, \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2, \tag{10}$$

(ii) The posterior covariances satisfies the bound $\lambda_{\max}\big(\Sigma_{j,n_j}(s)\big) \leq \min\Big\{\lambda_{\max}(\Sigma_j(s)), \frac{\varsigma_j^2}{n_j(s)}\Big\}$, where $\mathcal{U}_s(\beta_j) = \frac{n_j(s)\big(c(s,\beta_j)+q(s,\beta_j)/\varsigma_j^2\big)}{1+n_j(s)\big(c(s,\beta_j)+q(s,\beta_j)/\varsigma_j^2\big)}$, and $\lambda_{\max}(\cdot)$ denotes the maximum eigenvalue of its matrix argument.

Corollary 2 *Under the assumptions of Theorem 1, and assuming $\|\bar{\mathbf{x}}_j(s) - \bar{\mathbf{x}}_j\|_2 \to 0$ as $s \to \infty$:*

(i) If $\sigma_j(s) \to \infty$, then for any fixed β_j, as $s \to \infty$,

$$\underbrace{\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \to \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2,}_{\text{posterior mean moves away from the prior mean}} \qquad \underbrace{\|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \to 0.}_{\text{posterior mean moves toward the empirical cluster mean}} \tag{11}$$

(ii) If $\sigma_j(s) \to 0$, then for any fixed β_j, as $s \to \infty$,

$$\underbrace{\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \to 0,}_{\text{posterior mean moves toward the prior mean}} \qquad \underbrace{\|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \to \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2.}_{\text{posterior mean moves away from the empirical cluster mean}} \tag{12}$$

(iii) If $\sigma_j(s) \leq M < \infty$ for all $s \in \mathcal{T}$, then for fixed $s > 0$, $\mathcal{U}_s(\beta_j) \to 0$ as $\beta_j \to \infty$, implying

$$\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \to 0, \qquad \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \to \|\mathbf{m}_j - \bar{\mathbf{x}}_j(s)\|_2. \tag{13}$$

Periodic perturbation of centroid trajectories. Corollary 2(i) and (ii) reveal that increasing $\sigma_j(s)$ strengthens the pull of the posterior mean toward $\bar{x}_j(s)$ (ground-truth state), whereas letting $\sigma_j(s) \to 0$ drives it toward $m_j(s)$ (perturbed state). Motivated by this trade-off, we propose choosing $\sigma_j(s)$ to be periodic, inducing centroid trajectories that oscillate between ground-truth and perturbed states over time. To this end, we partition time into cycles of length T. Specifically, for a given sufficiently large $\sigma_j > 0$, we set $\sigma_j(s_0 + kT) = \sigma_j(s_0 + (k+1)T) = \sigma_j$, corresponding to the beginning and end of the k-th cycle, while $\sigma_j(s_1 + kT) = 0$. Here, s_0 is the cycle start time, $s_1 = s_0 + T/2$ is the cycle midpoint, and $k \in \mathbb{N}_0$ indexes the cycles.

Controlling perturbation intensity. For any fixed time s, Corollary 2(iii) shows that increasing β_j drives $\mathcal{U}_s(\beta_j)$ toward 0, drawing the posterior mean closer to the perturbed state. Since $\mathcal{U}_s(\beta_j) \in [0,1]$, it quantifies the diffusion-induced perturbation and serves as a utility-privacy trade-off parameter for centroid trajectories.

Building on the above settings, we now define diffusion privacy to protect centroid trajectory queries.

Definition 1 *For each $j \in [\![N]\!]$, let the centroid trajectory $\{\boldsymbol{\mu}_j(s) : s \in \mathcal{T}\}$ be modeled by the MR-OUP in equation 1 with periodic diffusion coefficient $\sigma_j(s)$. The trajectory is said to be $\mathcal{U}_s(\beta_j)$-diffusion private at time s if equation 9 and equation 10 hold.*

Here, privacy means trajectory protection through controlled perturbation, rather than a DP guarantee. Section 5 and Appendix A.4.3 examine periodic $\sigma_j(s)$ and their effects on centroid trajectories.

4.3 $(\epsilon(\alpha), \delta(\alpha))$-Differential privacy for weights

We next study how the Dirichlet weights in equation 5 control membership leakage through PP mixture weights, protecting group-proportion queries.

Theorem 2 *Let $\Delta_N = \{\boldsymbol{u} \in \mathbb{R}^N : u_i \geq 0,\ \sum_{i=1}^N u_i = 1\}$ and, for $\gamma_\alpha = [1-(1+\alpha)^{-q}]/N$, $q \in (0, 0.5)$, define $\Omega^{\pi}_{\gamma_\alpha} = \{\boldsymbol{\pi} \in \Delta_N : \pi_j \geq \gamma_\alpha,\ j \in [\![N]\!]\}$ and $\Omega^{w}_{\gamma_\alpha} = \{\boldsymbol{w} \in \Delta_N : w_j \geq \gamma_\alpha,\ j \in [\![N]\!]\}$. For every $s \in \mathcal{T}$, let $\mathbf{X}_{1:n}(s)$ and $\mathbf{X}'_{1:n}(s)$ be neighboring datasets differing in one entry such that $\boldsymbol{\pi}(\mathbf{X}_{1:n}(s)), \boldsymbol{\pi}(\mathbf{X}'_{1:n}(s)) \in \Omega^{\pi}_{\gamma_\alpha}$ and $\|\boldsymbol{\pi}(\mathbf{X}_{1:n}(s)) - \boldsymbol{\pi}(\mathbf{X}'_{1:n}(s))\|_1 \leq \frac{2}{n}$, where $\boldsymbol{\pi}(\mathbf{X}_{1:n}(s)) = (n_1(s)/n, \ldots, n_N(s)/n)$. Define*

$$\mathcal{M}(\boldsymbol{\pi}(\mathbf{X}_{1:n}(s))) := (w_1, \ldots, w_N) \mid z_{1:n}(s) \sim \mathrm{Dir}(\alpha/N + n\pi_1(s), \ldots, \alpha/N + n\pi_N(s))\,.$$

Let $(\Omega^{w}_{\gamma_\alpha})^c = \Delta_N \setminus \Omega^{w}_{\gamma_\alpha}$, and let ψ' denote the trigamma function. Then $\mathcal{M}$ is $(\epsilon(\alpha), \delta(\alpha))$-DP, where

$$\epsilon(\alpha) = 2n\psi'\left(\frac{\alpha}{N}\right) + \log\left(\frac{1-(N-1)\gamma_\alpha}{\gamma_\alpha}\right), \quad \delta(\alpha) = \sup_{\boldsymbol{\pi} \in \Omega^{\pi}_{\gamma_\alpha}} \Pr\big[\mathcal{M}(\boldsymbol{\pi}) \in (\Omega^{w}_{\gamma_\alpha})^c\big]\,.$$

Corollary 3 (Perfect Privacy Limit) *As $\alpha \to \infty$, Theorem 2 yields (i) $\epsilon(\alpha) \to 0$ and (ii) $\delta(\alpha) \to 0$.*

Remark 1 *Theorem 2 also controls membership leakage through expected weights in equation 6. For neighboring $X(s)$ and $X'(s)$, $\mathbb{E}[w_j \mid z] \leq e^{\epsilon(\alpha)}\mathbb{E}[w_j \mid z'] + \delta(\alpha)$, and conversely.*

5 Experimental results

We first evaluate the procedure on a toy example and then apply it to a real dataset. Throughout, we adaptively select N, the number of latent groups, for fixed α and $\rho = 0.05$; see Appendix A.4.1.

Toy Example: We consider a two-modal Gaussian population with $d = 1$ over $s \in [0, 20]$,

$$X(s) \sim 0.75\,\mathrm{N}(\mu_{\mathrm{true},1}(s), 0.16^2) + 0.25\,\mathrm{N}(\mu_{\mathrm{true},2}(s), 0.16^2),$$

where $\mu_{\mathrm{true},1}(s)$ and $\mu_{\mathrm{true},2}(s)$ are the true centroids of groups 1 and 2. Let $h_{[a,b]}(s) = \{\tanh((s-a)/0.08) - \tanh((s-b)/0.08)\}/2$, approximately one on $[a,b]$ and zero outside. We set

$$\begin{aligned}\mu_{\mathrm{true},1}(s) &= 4 + 0.045\sin(4s) + 0.025\sin(8s+0.4) - 1.25h_{[4,6]}(s)(1+0.02\sin(10s)) \\ &\quad + 0.85h_{[9,12]}(s)(1+0.025\sin(9s)), \\ \mu_{\mathrm{true},2}(s) &= 0.04\sin(4.3s+0.5) + 0.02\sin(8.5s) + 1.55h_{[9,12]}(s)(1+0.025\sin(9.5s+0.5)).\end{aligned}$$

Thus, group 1 jumps on $[4,6]$, while both groups jump on $[9,12]$, with a larger jump for group 2. We call these intervals *sensitive windows*. The corresponding true trajectories are shown in Figure 2. Our goal is to make it difficult to infer which group has no jump and which group has the larger jump. We do this by making the synthetic centroid trajectories have jumps of similar size. We first generate $\mathbf{x}_{1:50}(s)$ at 10^3 equally spaced time points over $[0,20]$. Next, for each interval, we model the centroid

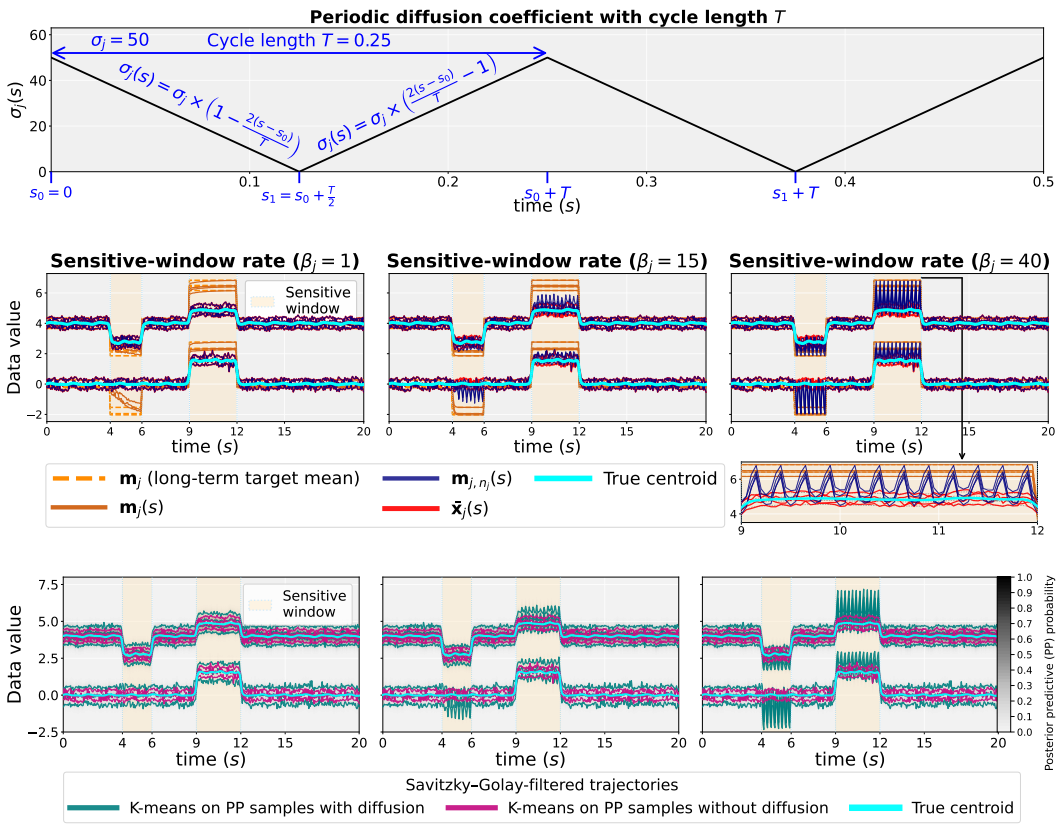

Figure 2: Toy example results. **Top:** Diffusion coefficient $\sigma_j(s)$. **Middle:** Posterior means for $\beta_j \in \{1, 15, 40\}$. **Bottom:** Filtered PP trajectories with and without diffusion versus true centroids.

trajectories $\mu_j(s)$ of all latent groups $j \in [\![N]\!]$ using the same MR-OUP, with different MR-OUPs across intervals. The MR-OUPs use a simple periodic $\sigma_j(s)$ shown in Figure 2 (top), with

$$\begin{aligned} m_j &= (\bar{x}_j(0) + r_{[0,4]})\mathbf{1}_{[0,4]}(s) + (\bar{x}_j(4) + r_{[4,6]})\mathbf{1}_{[4,6]}(s) + (\bar{x}_j(6) + r_{[6,9]})\mathbf{1}_{[6,9]}(s) \\ &\quad + (\bar{x}_j(9) + r_{[9,12]})\mathbf{1}_{[9,12]}(s) + (\bar{x}_j(12) + r_{[12,20]})\mathbf{1}_{[12,20]}(s), \\ \beta_j &= \beta_{[0,4]}\mathbf{1}_{[0,4]}(s) + \beta_{[4,6]}\mathbf{1}_{[4,6]}(s) + \beta_{[6,9]}\mathbf{1}_{[6,9]}(s) + \beta_{[9,12]}\mathbf{1}_{[9,12]}(s) + \beta_{[12,20]}\mathbf{1}_{[12,20]}(s). \end{aligned}$$

Evaluating $\mathcal{U}_s(\beta_j)$-diffusion. In sensitive windows, we set jump sizes $r_{[4,6]} = -1.8$, $r_{[9,12]} = 2.5$, and $\beta_{[4,6]} = \beta_{[9,12]} \in \{1, 15, 40\}$. Outside these windows, we set $r_{[0,4]} = r_{[6,9]} = r_{[12,20]} = 0$ and $\beta_{[0,4]} = \beta_{[6,9]} = \beta_{[12,20]} = 1$. To examine the behavior of $\mu_j(s)$, we focus on its posterior mean $m_{j,n_j}(s)$. For each interval $[a, b]$ with $s_0 = a$, we apply Lemma 1 with $m_j(a) = \bar{x}_j(a)$. For $\Sigma_j \sim \text{IW}(\Lambda_0, \nu_0)$, following Murphy (2012), we set $\nu_0 = d + 2$ and $\Lambda_0 = \frac{1}{n_j}\{S_{\bar{x}_j(0)}\mathbf{1}_{[0,4]}(s) + S_{\bar{x}_j(4)}\mathbf{1}_{[4,6]}(s) + S_{\bar{x}_j(6)}\mathbf{1}_{[6,9]}(s) + S_{\bar{x}_j(9)}\mathbf{1}_{[9,12]}(s) + S_{\bar{x}_j(12)}\mathbf{1}_{[12,20]}(s)\}$, where $S_{\bar{x}_j(a)}$ is the sample covariance. We then compute $m_j(s)$ and $\Sigma_j(s)$ from Lemma 1 and use them to obtain $m_{j,n_j}(s)$ from the fitted DDirPMM, shown in Figure 2 (middle). As β_j increases in sensitive windows, $m_{j,n_j}(s)$ fluctuates more between $\bar{x}_j(s)$ and $m_j(s)$, consistent with Corollary 2(iii). At $\beta_j = 40$, the posterior mean trajectories reach similar jump sizes, making the underlying jumps harder to distinguish. The zoomed-in panel also shows how the periodic pattern of $\sigma_j(s)$ shapes these fluctuations. Appendix A.4.3 considers a more complex $\sigma_j(s)$ to illustrate its flexibility in generating more complex trajectories.

We next test whether a third party can recover centroid trajectories from synthetic data using repeated clustering and filtering. To isolate the effect of diffusion, we fix $\alpha = 10^{-5}$, minimizing the DP effect of mixture weights on the synthetic samples. At each time s, we draw 50 PP samples and apply K-means with $K = N$ over 100 repetitions. The estimated centroids are averaged and smoothed using a Savitzky-Golay filter (Savitzky & Golay, 1964), as shown in Figure 8 (bottom). Consistent with the posterior means, larger β_j makes the synthetic centroid trajectories vary more between the true and perturbed states as $\mathcal{U}_s(\beta_j)$ varies more over $[0, 1]$ (Figure 3). Figure 3 also reports the

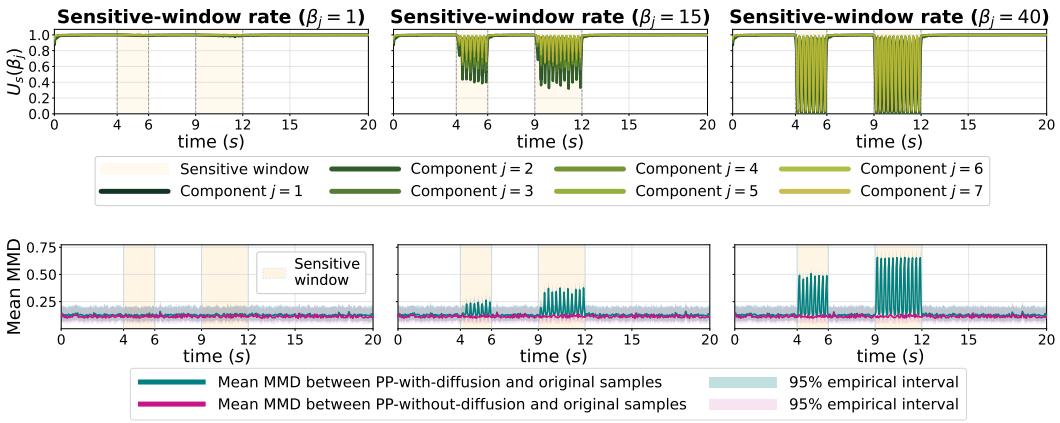

Figure 3: Privacy-utility trade-off induced by diffusion trajectories in the toy example. **Top:** $\mathcal{U}_s(\beta_j)$ across s for different β_j. **Bottom:** MMD between PP and original samples.

maximum mean discrepancy (MMD) between the synthetic and original samples. The MMD varies within sensitive windows rather than remaining high throughout, avoiding persistent utility loss.

To our knowledge, no existing method directly protects dynamic centroid trajectories, so we compare against a model without diffusion. We replace the diffusion model in Lemma 1 with $\mu_j(s) \mid \Sigma_j \sim \mathrm{N}(m_j, \Sigma_j)$. The filtered centroid trajectories without diffusion are shown in Figure 2 (bottom). Since $m_{j,n_j} = \{m_j + n_j \bar{x}_j\}/(1 + n_j)$ gives more weight to $\bar{x}_j$ for $n_j > 1$, the trajectories remain close to the true centroids. Thus, the static model provides limited protection against trajectory reconstruction.

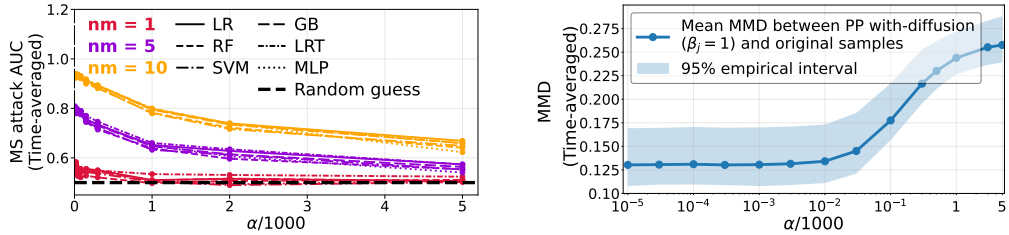

Figure 4: Privacy-utility trade-off induced by Dirichlet weights in the toy example. **Left:** Time-averaged membership attack AUC. **Right:** Time-averaged MMD between PP and original samples.

Evaluating $(\epsilon(\alpha), \delta(\alpha))$-DP. To isolate the effect of DP, we fix $\beta_j = 1$ in both sensitive and non-sensitive windows and vary α. See Appendix A.4.2 for $\epsilon(\alpha), \delta(\alpha)$ versus α. We aim to evaluate how DP on mixture weights limits membership (MS) leakage through changes in group proportions, as established in Theorem 2. We consider two datasets, $\mathbf{X}(s)$ and $\mathbf{X}'(s)$, that differ in the membership of nm observations. $\mathbf{X}'(s)$ is obtained from $\mathbf{X}(s)$ by replacing nm observations with observations drawn from group 2 of the toy population. Standard record-level DP corresponds to the neighboring case $nm = 1$. We also consider $nm \in \{5, 10\}$ to assess protection under larger group-level changes.

We assess this protection using five attackers: Logistic Regression (LR), Random Forest (RF), Support Vector Machine (SVM), Multi-Layer Perceptron (MLP), and Gradient Boosting (GB), together with a likelihood-ratio test (LRT) attacker. The PP mixture weights in equation 6 corresponding to $\mathbf{X}(s)$ and $\mathbf{X}'(s)$ are used to train and evaluate the attackers, which infer which dataset the weights correspond to; see Appendix A.4.4. We evaluate performance using the area under the ROC curve (AUC) (Fawcett, 2006), where 0.5 corresponds to random guessing. We compute AUC at each $s \in [0, 20]$ and average it over time. Figure 4 (left) shows that the time-averaged AUC approaches 0.5 as α increases, indicating that the attackers become less able to infer whether the mixture weights come from $\mathbf{X}(s)$ or $\mathbf{X}'(s)$. Larger nm, however, leads to stronger attacks.

To evaluate utility, we compute the MMD between the PP and original samples at each time point $s \in [0, 20]$ and then average it over time. Figure 4 (right) shows that the time-averaged MMD remains nearly unchanged for small α but increases as α becomes larger, indicating greater utility loss. Together with the results in Figure 4 (left), this shows the privacy-utility trade-off: larger α provides stronger weight privacy but can reduce utility.

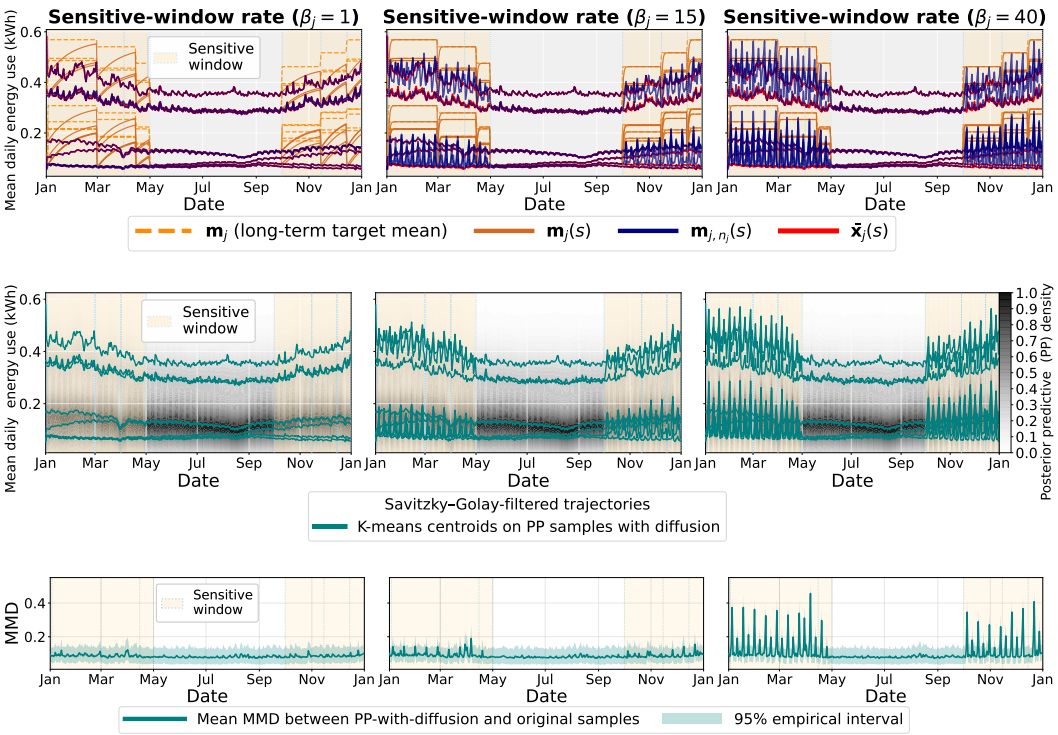

Figure 5: GoiEner Smart Meters results. **Top:** Posterior means. **Middle:** Filtered PP trajectories. **Bottom:** MMD between PP and original samples.

GoiEner Smart Meters Data. We use daily electricity consumption from 7,975 customers across peninsular Spain in 2018 (Quesada et al., 2024), using each customer's average hourly consumption per day. To display the underlying centroid trajectories, we first minimize diffusion by setting $m_j = \bar{x}_j(\text{Jan1})\mathbf{1}_{[\text{Jan1,Dec31}]}$ and $\beta_j = \mathbf{1}_{[\text{Jan1,Dec31}]}$, then obtain $m_{j,n_j}(s)$ from the fitted DDirPMM. Figure 5 (top-left) shows distinct group patterns, with consumption increasing from October to January and decreasing from January to May (sensitive windows). These patterns may reflect greater electricity use for heating during colder months, potentially exposing some groups to targeted promotions from electric-appliance companies and creating unequal opportunities. To protect these patterns, we induce a common trend across latent groups in the sensitive windows. For smooth transitions, we divide each window into three intervals and use different MR-OUPs across intervals, with

$$\begin{aligned} m_j = {} & (\bar{x}_j(\text{Jan1}) + r_{[\text{Jan1,Mar1}]})\mathbf{1}_{[\text{Jan1,Mar1}]}(s) + (\bar{x}_j(\text{Mar1}) + r_{[\text{Mar1,Apr1}]})\mathbf{1}_{[\text{Mar1,Apr1}]}(s) \\ & + (\bar{x}_j(\text{Apr1}) + r_{[\text{Apr1,May1}]})\mathbf{1}_{[\text{Apr1,May1}]}(s) + (\bar{x}_j(\text{May1}) + r_{[\text{May1,Oct1}]})\mathbf{1}_{[\text{May1,Oct1}]}(s) \\ & + (\bar{x}_j(\text{Oct1}) + r_{[\text{Oct1,Oct15}]})\mathbf{1}_{[\text{Oct1,Oct15}]}(s) + (\bar{x}_j(\text{Oct15}) + r_{[\text{Oct15,Dec15}]})\mathbf{1}_{[\text{Oct15,Dec15}]}(s) \\ & + (\bar{x}_j(\text{Dec15}) + r_{[\text{Dec15,Dec31}]})\mathbf{1}_{[\text{Dec15,Dec31}]}(s). \end{aligned}$$

We set the jump sizes to $r_{[\text{Jan1,Mar1}]} = 0.09$, $r_{[\text{Mar1,Apr1}]} = 0.06$, $r_{[\text{Apr1,May1}]} = 0.05$, $r_{[\text{May1,Oct1}]} = 0$, $r_{[\text{Oct1,Oct15}]} = 0.08$, $r_{[\text{Oct15,Dec15}]} = 0.09$, and $r_{[\text{Dec15,Dec31}]} = 0.12$. We use the same β_j settings as in the toy example and the complex periodic $\sigma_j(s)$ in Appendix A.4.3. Figure 5 (top) shows the designed m_j and corresponding $m_{j,n_j}(s)$. As β_j increases in the sensitive windows, the distinct patterns become harder to infer. The middle shows Savitzky-Golay-smoothed synthetic centroid trajectories, with a similar interpretation. The bottom shows the effect on utility. Throughout, we set $\alpha = 500$ to impose only a weak $(26.2, 0)$-DP effect on the weights.

6 Concluding Remarks

We identified a new privacy concern in dynamic clustering. To address it, we proposed a synthetic-data mechanism using diffusion to protect centroid trajectories from reconstruction, while providing DP protection for mixture weights. Theoretical and empirical results support both mechanisms. This work also opens new privacy questions for dynamic mixture models, including evolving mixture weights in interacting LLMs (Wang et al., 2025) and formal DP guarantees for centroid trajectories. These emerging directions provide promising opportunities for future research; see Appendix A.7.

AI USE DISCLOSURE

In this work, we used generative AI tools for proofreading and improving the grammatical clarity and readability of the manuscript. We have reviewed all AI-assisted work and take responsibility for the final content of this work.

ETHICS STATEMENT

This work uses synthetic data and the publicly available GoiEner Smart Meters dataset. No new human-subject data were collected.

REPRODUCIBILITY STATEMENT

Details supporting the reproducibility of the theoretical and experimental results are provided throughout the main text and appendix, including proofs, experimental settings, implementation details, and algorithms.

REFERENCES

Rebecca Adam, Kun Qian, and Robert Brehm. Electric vehicle user behavior prediction using gaussian mixture models and soft information. In *2021 IEEE PES Innovative Smart Grid Technologies-Asia (ISGT Asia)*, pp. 1–5. IEEE, 2021.

Jozsef Bakosi and JR Ristorcelli. A stochastic diffusion process for lochner's generalized dirichlet distribution. *Journal of Mathematical Physics*, 54(10), 2013.

Mario Beraha, Stefano Favaro, and Vinayak Rao. Mcmc for bayesian nonparametric mixture modeling under differential privacy. *Journal of Computational and Graphical Statistics*, 34(3):837–847, 2025.

David M Blei and Michael I Jordan. Variational inference for dirichlet process mixtures. *Bayesian Analysis*, 1(1):121–144, 2006.

Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.

Gianfranco Chicco. Overview and performance assessment of the clustering methods for electrical load pattern grouping. *Energy*, 42(1):68–80, 2012.

Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20(3):273–297, 1995.

David R Cox. The regression analysis of binary sequences. *Journal of the Royal Statistical Society: Series B (Methodological)*, 21(1):238–238, 1959.

Giuseppe Da Prato and Alessandra Lunardi. Ornstein–uhlenbeck operators with time periodic coefficients. *Journal of Evolution Equations*, 7(4):587–614, 2007.

Yi Dai, Shuo Liu, Hao Li, Qi Chen, Xiaochen Liu, Xiaohua Liu, and Tao Zhang. Pattern extraction and structured characterization for electricity consumption profiles in different types of buildings. *Energy and Buildings*, 336:115598, 2025.

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

Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.

Thomas S Ferguson. A bayesian analysis of some nonparametric problems. *The annals of statistics*, pp. 209–230, 1973.

Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pp. 1189–1232, 2001.

Quentin Gemine, Bertrand Cornélusse, Mevludin Glavic, Raphael Fonteneau, and Damien Ernst. A gaussian mixture approach to model stochastic processes in power systems. In *2016 Power Systems Computation Conference (PSCC)*, pp. 1–7. IEEE, 2016.

Zhitao Guan, Zefang Lv, Xianwen Sun, Longfei Wu, Jun Wu, Xiaojiang Du, and Mohsen Guizani. A differentially private big data nonparametric bayesian clustering algorithm in smart grid. *IEEE Transactions on Network Science and Engineering*, 7(4):2631–2641, 2020.

Antti Honkela and Laila Melkas. Gaussian processes with differential privacy. *arXiv preprint arXiv:2106.00474*, 2021.

Roger A Horn and Charles R Johnson. *Matrix Analysis*. Cambridge University Press, 1990.

Hemant Ishwaran and Mahmoud Zarepour. Exact and approximate sum representations for the Dirichlet process. *Canadian Journal of Statistics*, 30(2):269–283, 2002.

Paul A. Jenkins and Dario Spanò. Exact simulation of the Wright–Fisher diffusion. *The Annals of Applied Probability*, 27(3):1478–1509, 2017. doi: 10.1214/16-AAP1236.

Michael I Jordan. Dirichlet processes, chinese restaurant processes and all that. In *Tutorial presentation at the NIPS Conference*, 2005.

Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

Jiaqi Li, Huichao Kong, Huafeng Su, Yinping Dai, Zixian Zeng, Jiancong Lu, and Jinxing Zhong. Real-time abnormal alarm method for distribution network equipment based on gaussian mixture model. In *2025 10th Asia Conference on Power and Electrical Engineering (ACPEE)*, pp. 1023–1028. IEEE, 2025.

Qiongxiu Li, Richard Heusdens, and Mads Græsbøll Christensen. Privacy-preserving distributed optimization via subspace perturbation: A general framework. *IEEE Transactions on Signal Processing*, 68:5983–5996, 2020.

Qiongxiu Li, Jaron Skovsted Gundersen, Katrine Tjell, Rafal Wisniewski, and Mads Græsbøll Christensen. Privacy-preserving distributed expectation maximization for gaussian mixture model using subspace perturbation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4263–4267. IEEE, 2022.

Hang Liu, Anna Scaglione, and Sean Peisert. Differentially private distribution release of gaussian mixture models via kl-divergence minimization. *arXiv preprint arXiv:2506.03467*, 2025.

Carl D Meyer. *Matrix Analysis and Applied Linear Algebra*, volume 71. SIAM, 2000.

Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.

Iing Muttakhiroh and Thomas Fevens. Enhancing memory recall in llms with gauss-tin: A hybrid instructional and gaussian replay approach. *arXiv preprint arXiv:2508.09510*, 2025.

Soroush Najafi and Hanif Livani. Robust day-ahead voltage support and building demand response scheduling under gaussian mixture model uncertainty. *IEEE Transactions on Industry Applications*, 2025.

Jerzy Neyman and Egon Sharpe Pearson. On the problem of the most efficient tests of statistical hypotheses. In *Breakthroughs in Statistics: Foundations and Basic Theory*, pp. 73–108. Springer, 1992.

Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.

Carlos Quesada, Leire Astigarraga, Chris Merveille, and Cruz E Borges. An electricity smart meter dataset of spanish households: insights into consumption patterns. *Scientific Data*, 11(1):59, 2024.

Fernando A Quintana, Peter Müller, Alejandro Jara, and Steven N MacEachern. The dependent dirichlet process and related models. *Statistical Science*, 37(1):24–41, 2022.

Nikhil Ravi, Anna Scaglione, Sachin Kadam, Reinhard Gentz, Sean Peisert, Brent Lunghino, Emmanuel Levijarvi, and Aram Shumavon. Differentially private k-means clustering applied to meter data analysis and synthesis. *IEEE transactions on smart grid*, 13(6):4801–4814, 2022.

Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.

David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

Abraham Savitzky and Marcel JE Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964.

Kangwon Seo, Hyeong Suk Na, Wonjae Lee, Cheng-Bang Chen, Sang Jin Kweon, Long Zhao, and Soundar Kumara. Clustering electricity consumption patterns using functional data analysis. *Sustainable Energy, Grids and Networks*, 43:101742, 2025.

Jayaram Sethuraman. A constructive definition of Dirichlet priors. *Statistica Sinica*, pp. 639–650, 1994.

Ankita Sharma, Amrit Lal Sangal, and Kailash Chand Sharma. Privacy-preserving prosumer profiling using smart meter data. *Sustainable Energy, Grids and Networks*, 43:101784, 2025.

Miao Su, Yucan Guo, Zhongni Hou, Long Bai, Zixuan Li, Yufei Zhang, Guojun Yin, Wei Lin, Xiaolong Jin, Jiafeng Guo, et al. Beyond dialogue time: Temporal semantic memory for personalized llm agents. *arXiv preprint arXiv:2601.07468*, 2026.

Xiaosu Tong, Bowei Xi, Murat Kantarcioglu, and Ali Inan. Gaussian mixture models for classification and hypothesis tests under differential privacy. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pp. 123–141. Springer, 2017.

Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Cinzia Viroli and Geoffrey J McLachlan. Deep gaussian mixture models. *Statistics and Computing*, 29(1):43–51, 2019.

Chong Wang and David M Blei. Variational inference in nonconjugate models. *Journal of Machine Learning Research*, 2013.

Edward L Wang, Tianyu Wang, Hayden Helm, Avanti Athreya, Vince Lyzinski, and Carey E Priebe. Gaussian mixture models as a proxy for interacting language models. *arXiv preprint arXiv:2506.00077*, 2025.

Weisan Wu. A study on pufferfish privacy algorithm based on gaussian mixture models. *Scientific Reports*, 15(1):1015, 2025.

Yuncheng Wu, Yao Wu, Hui Peng, Juru Zeng, Hong Chen, and Cuiping Li. Differentially private density estimation via gaussian mixtures model. In *2016 IEEE/ACM 24th International Symposium on Quality of Service (IWQoS)*, pp. 1–6. IEEE, 2016.

Wenpeng Xing, Zhipeng Chen, Changting Lin, and Meng Han. Hgmf: A hierarchical gaussian mixture framework for scalable tool invocation within the model context protocol. *arXiv preprint arXiv:2508.07602*, 2025.

Liang Xiong, Barnabás Póczos, Jeff Schneider, Andrew Connolly, and Jake VanderPlas. Hierarchical probabilistic models for group anomaly detection. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 789–797. JMLR Workshop and Conference Proceedings, 2011.

Derong Xu, Yi Wen, Pengyue Jia, Yingyi Zhang, Yichao Wang, Huifeng Guo, Ruiming Tang, Xiangyu Zhao, Enhong Chen, Tong Xu, et al. From single to multi-granularity: Toward long-term memory association and selection of conversational agents. *arXiv preprint arXiv:2505.19549*, 2025.

Mahmoud Zarepour and Luai Al-Labadi. On a rapid simulation of the Dirichlet process. *Statistics & Probability Letters*, 82(5):916–924, 2012.

Yixiang Zhang, Xinhao Deng, Zhongyi Gu, Yihao Chen, Ke Xu, Qi Li, and Jianping Wu. Exposing llm user privacy via traffic fingerprint analysis: A study of privacy risks in llm agent interactions. *arXiv preprint arXiv:2510.07176*, 2025.

A Appendix

A.1 Notation

Unless stated otherwise, bold symbols denote vectors, uppercase $\mathbf{X}$ and Z denote random variables, and lowercase $\mathbf{x}$ and z denote their realizations. For a sequence $a_1, a_2, \ldots$, let $a_{1:m} = (a_1, \ldots, a_m)$.

Dimensions, Data, and Indexing

n, d, d'	Sample size, data dimension, output dimension of the generic DP mechanism, respectively.
K	Number of clusters used by K-means or the membership-attack reconstruction.
$\mathcal{T}$	Time domain.
$\mathbf{X}_i$, $\mathbf{x}_i$, P	Generic random observation, its realization, and the data-generating distribution.
$\mathbf{X}_i(s)$, $\mathbf{x}_i(s)$, P_s	ith random variable at time s, its observed realization, and the time-s marginal distribution, respectively.
$\mathbf{X}_{1:n}(s)$, $\mathbf{x}_{1:n}(s)$	Random dataset at time s and its observed realization, respectively.
$Z_i(s)$, $z_i(s)$, $z_{1:n}(s)$	Random component allocation, its realization, and the vector of realized allocations at time s.
$I_j(s)$, $n_j(s)$	Set of indices assigned to component j and its size, $n_j(s)$, respectively.
$\mathbf{X}_{I_j}(s)$, $\mathbf{x}_{I_j}(s)$	Random and observed collections assigned to component j at time s.
$\bar{\mathbf{x}}_j(s)$	Empirical mean of component j at time s.
$\mathbf{0}$, I_d	Zero vector, and the $d \times d$ identity matrix.

Sets and Domains

$(\mathfrak{X}, \mathcal{A})$	Measurable space, with sample space $\mathfrak{X}$ and σ-algebra $\mathcal{A}$.
$\mathbb{R}$, $\mathbb{R}^d$, $\mathbb{R}^{d\times d}$	Real numbers, d-dimensional real vector space, and space of real $d \times d$ matrices, respectively.
$\mathbb{N}_0$	Nonnegative integers.
$\lvert A\rvert$	Cardinality of set A, i.e., the number of elements in A.
$[\![m]\!]$	The set of all integers between 1 and m, i.e., $\{1, \ldots, m\}$, the manuscript uses $m \in \{n, N, d\}$.
$[a, b]$	The real interval including a and b
$(a, b]$	The real interval excluding a but including b
$\mathbb{A}\backslash\mathbb{B}$	Set subtraction, i.e., the set containing the elements of $\mathbb{A}$ that are not in $\mathbb{B}$
Δ_N	Probability simplex $\{\boldsymbol{u} \in \mathbb{R}^N : u_i \geq 0,\ \sum_{i=1}^N u_i = 1\}$.

$\Omega^w_{\gamma_\alpha}$, $(\Omega^w_{\gamma_\alpha})^c$	Subset of Δ_N with each weight $w_j \geq \gamma_\alpha$, and its complement, respectively.
$\Omega^\pi_{\gamma_\alpha}$, $(\Omega^\pi_{\gamma_\alpha})^c$	Subset of Δ_N with each empirical group proportion $\pi_j(s) \geq \gamma_\alpha$, and its complement, respectively.

Probability, Distributions, and Operators

$\Pr[\cdot]$, $\mathbb{E}[\cdot]$, $\mathrm{Var}[\cdot]$	Probability, expectation, and variance under the governing distribution.
$\mathrm{N}(\boldsymbol{\mu}, \Sigma)$	Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance Σ; the same notation is used for scalar Gaussian distributions.
$\mathrm{Dir}(a_1, \ldots, a_N)$	Dirichlet distribution with parameter vector $(a_1, \ldots, a_N)$.
$\mathrm{DP}(\alpha, \mathring{G})$	Dirichlet process with concentration parameter α and base measure $\mathring{G}$; "DP" without arguments denotes differential privacy.
$\mathrm{IW}(\Lambda, \nu)$	Inverse-Wishart distribution with scale matrix Λ and degrees of freedom ν.
$\mathrm{Gamma}(a, b)$	Gamma distribution with shape a and rate b.
$\mathrm{D}_{\mathrm{KL}}(q\|p)$	Kullback–Leibler divergence from density q to p.
$\xrightarrow{p}$, $\xrightarrow{d}$	Convergence in probability and convergence in distribution.
$X \perp\!\!\!\perp Y$	Random variables X and Y are independent.
$\|\mathbf{v}\|_1$, $\|\mathbf{v}\|_2$	ℓ_1 and Euclidean norms of vector $\mathbf{v}$, respectively.
$\|\mathrm{A}\|_{\mathrm{op}}$	Induced operator 2-norm of matrix A.
$\mathrm{A}^\top$, $\det(\mathrm{A})$, $\mathrm{tr}(\mathrm{A})$	Transpose, determinant, and trace.
$\mathrm{diag}(\mathbf{v})$	Diagonal matrix with vector $\mathbf{v}$ as its diagonal.
$\mathrm{card}(S)$	Cardinality, i.e., the size of set S.
$\lambda_i(\mathrm{A})$, $\lambda_{\max}(\mathrm{A})$	The ith eigenvalue and the largest eigenvalue of matrix A.
$\Gamma(x)$, $\psi(x)$, $\psi'(x)$	Gamma, digamma, and trigamma functions, respectively.
∇_x	Gradient with respect to x.
$\mathrm{softplus}(x)$	Softplus function, $\log(1 + e^x)$.
$\mathsf{D}_x(\cdot)$	Dirac probability measure (unit point mass) at atom x.
$\mathbf{1}_{[a,b]}(s)$	Indicator function taking value 1 if $s \in [a, b]$ and 0 otherwise.
$\mathrm{A} \succeq \mathrm{B}$	$\mathrm{A} - \mathrm{B}$ is positive semidefinite.

Dirichlet-Process Mixture Model

G	A random probability measure follows from $\mathrm{DP}(\alpha, \mathring{G})$.

α	DirP concentration parameter, taking positive real values.
$\mathring{G}$	Fixed probability measure serving as the DirP base measure.
$\mathring{G}'$, $\mathring{G}''$	Conditional and marginal decompositions of $\mathring{G}$, respectively.
$\mathring{g}$, $\mathring{g}'$, $\mathring{g}''$	Density functions corresponding to $\mathring{G}$, $\mathring{G}'$, $\mathring{G}''$, respectively.
$\kappa(\cdot \mid \boldsymbol{\theta})$	Parametric likelihood kernel of the DirPMM, parameterized by $\boldsymbol{\theta}$.
$p(\cdot \mid G)$	DirPMM prior density.
$p_N(\cdot \mid G)$	finite DirPMM prior density.
N	Truncation level in $p_N(\cdot \mid G)$.
$\omega_{1:N}$	Random weights in the finite DirPMM.
$g_j(\cdot \mid \mathbf{x}_{I_j}(s))$	Posterior predictive density of the j-th component.
$\dot{g}(\cdot)$	Prior predictive density associated with a new cluster.
Λ_0, ν_0	Scale matrix and degrees-of-freedom hyperparameter of the inverse-Wishart prior, respectively.

Diffusion Process and Posterior Quantities

$\boldsymbol{\mu}_j(s)$	Centroid of the jth component in the DDirPMM, modeled by an MR-OUP.
$\mathbf{m}_j$	Long-term equilibrium or data-owner-selected target means of the MR-OUP $\boldsymbol{\mu}_j(s)$.
$\mathbf{m}_j(s)$	Prior mean of $\boldsymbol{\mu}_j(s)$ induced by the MR-OUP at time s.
β_j	Mean-reversion rate of the MR-OUP $\boldsymbol{\mu}_j(s)$.
$\sigma_j(s)$	Time-varying diffusion coefficient of the MR-OUP $\boldsymbol{\mu}_j(s)$.
σ_j	Limiting value of $\sigma_j(s)$.
$\mathbf{W}(s)$	Standard d-dimensional Wiener process.
$\Sigma_j(s)$	Covariance matrix of the MR-OUP $\boldsymbol{\mu}_j(s)$ at time s.
Σ_j	Initial covariance matrix of the MR-OUP $\boldsymbol{\mu}_j(s)$ at time s_0, i.e., $\Sigma_j = \Sigma_j(s_0)$.
$\mathbf{m}_{j,n_j}(s)$	Posterior mean of $\boldsymbol{\mu}_j(s)$ at time s given the observations $\mathbf{x}_{I_j}(s)$.
$\Sigma_{j,n_j}(s)$	Posterior covariance of $\boldsymbol{\mu}_j(s)$ at time s given the observations $\mathbf{x}_{I_j}(s)$.
ς_j	Isotropic observation standard deviation in $\Sigma_j = \varsigma_j^2 \mathrm{I}_d$.
$\mathcal{U}_s(\beta_j)$	Diffusion-privacy coefficient, equal to $\frac{n_j(s)[c(s,\beta_j)+q(s,\beta_j)/\varsigma_j^2]}{1+n_j(s)[c(s,\beta_j)+q(s,\beta_j)/\varsigma_j^2]}$ when $\Sigma_j = \varsigma_j^2 \mathrm{I}_d$.
T, s_0, s_1, k	Cycle length, cycle start time, cycle midpoint $s_1 = s_0 + T/2$, and cycle index $k \in \mathbb{N}_0$, respectively.

Differential Privacy for Mixture Weights

$\mathbf{X}_{1:n}(s)$, $\mathbf{X}'_{1:n}(s)$	Neighboring datasets in the definition of differential privacy.
$\boldsymbol{\pi}(\mathbf{X}_{1:n}(s))$, $\pi_j(s)$	Empirical component-proportion vector and its component $\pi_j(s) = n_j(s)/n$.
$\mathcal{M}(\cdot)$	Randomized mechanism in DP.
$\epsilon(\alpha)$, $\delta(\alpha)$	Privacy parameters established for the Dirichlet weight mechanism as functions of α.
γ_α	Coordinate lower bounds defining Ω_{γ_α}.

Variational-Inference and Algorithmic Notation

$q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\vartheta})$, q_j	Mean-field variational approximation and its optimized component-j approximation, respectively.
$\mathcal{V}$	Variational parameter space in Appendix A.5.
$\boldsymbol{\vartheta}$, $\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}$, $\boldsymbol{\vartheta}_{\Sigma_j}$, $\boldsymbol{\vartheta}_j^*$	Full variational parameter vector, flow-specific parameter blocks for $\boldsymbol{\mu}_j$ and Σ_j, and optimized parameter vector, respectively.
$f_{\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}}$, $h_{\boldsymbol{\vartheta}_{\Sigma_j}}$	Normalizing-flow maps for the centroid and covariance representations, respectively.
J_f, J_h	Jacobian matrices of the corresponding normalizing-flow maps, respectively.
n_{MC}	Number of Monte Carlo draws used to estimate the ELBO.
$\boldsymbol{\epsilon}_{\boldsymbol{\mu}_j}^{(\ell)}$, $\boldsymbol{\epsilon}_{\Sigma_j}^{(\ell)}$	Standard Gaussian base-noise vectors for the centroid and covariance flows, respectively, at Monte Carlo draw ℓ.
$\mathcal{I}(k, m)$	Fixed map from a lower-triangular index pair (k, m) to its position in $\mathbf{r}_j^{(\ell)}$.
$\widehat{\mathcal{L}}_j$	Monte Carlo estimate of the component-j evidence lower bound (ELBO).

Acronyms

AUC	Area under the receiver operating characteristic curve.
BNP	Bayesian Nonparametric
CRP	Chinese restaurant process.
DirP	Dirichlet Process
DirPMM	Dirichlet process mixture model.
DDirPMM	Dynamic Dirichlet process mixture model.
DP	Differential privacy.
ELBO	Evidence lower bound.
GB	Gradient boosting.
GMM	Gaussian mixture model.

GP	Gaussian process.
i.i.d.	Independent and identically distributed.
IW	Inverse-Wishart.
KL	Kullback-Leibler Divergence
LLM	Large Language Model
LR	Logistic Regression classifier
LRT	Likelihood-ratio test.
MC	Monte Carlo.
MLP	Multi-layer perceptron.
MMD	Maximum mean discrepancy.
MR-OUP	Mean-reverting Ornstein–Uhlenbeck process.
MS	Membership.
MSE	Mean squared error.
PP	Posterior predictive.
RF	Random forest.
ROC	Receiver operating characteristic.
SDE	Stochastic differential equation.
SVM	Support vector machine.
VI	Variational inference.

A.2 Theoretical Proofs

Proof of Lemma 1

Recall that the Itô SDE in equation 1 is given by

$$d\boldsymbol{\mu}_i(s) = \beta_i\big(\mathbf{m}_i - \boldsymbol{\mu}_i(s)\big)\, ds + \sigma_i(s)\, d\mathbf{W}(s).$$

Define the centered process

$$\boldsymbol{\delta}_i(s) := \boldsymbol{\mu}_i(s) - \mathbf{m}_i. \tag{14}$$

Then

$$d\boldsymbol{\delta}_i(s) = -\beta_i \boldsymbol{\delta}_i(s)\, ds + \sigma_i(s)\, d\mathbf{W}(s).$$

Introduce the integrating factor $e^{\beta_i s}$ and define

$$\widetilde{\boldsymbol{\delta}}_i(s) := e^{\beta_i s}\boldsymbol{\delta}_i(s). \tag{15}$$

Applying Itô's formula gives

$$\begin{aligned} d\widetilde{\boldsymbol{\delta}}_i(s) &= e^{\beta_i s}\, d\boldsymbol{\delta}_i(s) + \beta_i e^{\beta_i s}\boldsymbol{\delta}_i(s)\, ds \\ &= e^{\beta_i s}\sigma_i(s)\, d\mathbf{W}(s). \end{aligned}$$

Integrating over $[s_0, s]$ gives

$$\widetilde{\boldsymbol{\delta}}_i(s) = \widetilde{\boldsymbol{\delta}}_i(s_0) + \int_{s_0}^{s} e^{\beta_i u}\sigma_i(u)\, d\mathbf{W}(u).$$

Hence,

$$\boldsymbol{\delta}_i(s) = e^{-\beta_i(s-s_0)}\boldsymbol{\delta}_i(s_0) + \int_{s_0}^{s} e^{-\beta_i(s-u)}\sigma_i(u)\, d\mathbf{W}(u). \tag{16}$$

Substituting equation 14 into equation 16, we obtain

$$\boldsymbol{\mu}_i(s) = \mathbf{m}_i + e^{-\beta_i(s-s_0)}\big(\boldsymbol{\mu}_i(s_0) - \mathbf{m}_i\big) + \int_{s_0}^{s} e^{-\beta_i(s-u)}\sigma_i(u)\, d\mathbf{W}(u). \tag{17}$$

By Oksendal (2013, Section 3), the Wiener integral

$$\int_{s_0}^{s} e^{-\beta_i(s-u)}\sigma_i(u)\, d\mathbf{W}(u)$$

is normally distributed with

$$\mathbb{E}\left[\int_{s_0}^{s} e^{-\beta_i(s-u)}\sigma_i(u)\, d\mathbf{W}(u)\right] = 0, \tag{18a}$$

$$\mathrm{Var}\left[\int_{s_0}^{s} e^{-\beta_i(s-u)}\sigma_i(u)\, d\mathbf{W}(u)\right] = \mathrm{I}_d \int_{s_0}^{s} e^{-2\beta_i(s-u)}\sigma_i^2(u)\, du. \tag{18b}$$

Since $\boldsymbol{\mu}_i(s_0) \sim \mathrm{N}(\mathbf{m}_i(s_0), \Sigma_i)$ and is independent of the Wiener increments, equation 17 is Gaussian with

$$\mathbb{E}\big(\boldsymbol{\mu}_i(s)\big) = \mathbf{m}_i + e^{-\beta_i(s-s_0)}\big(\mathbf{m}_i(s_0) - \mathbf{m}_i\big) := \mathbf{m}_i(s), \tag{19a}$$

$$\mathrm{Var}\big(\boldsymbol{\mu}_i(s)\big) = e^{-2\beta_i(s-s_0)}\Sigma_i + \mathrm{I}_d \int_{s_0}^{s} e^{-2\beta_i(s-u)}\sigma_i^2(u)\, du := \Sigma_i(s). \tag{19b}$$

Thus, retaining the notation in the main text,

$$c(s,\beta_i) = e^{-2\beta_i(s-s_0)}, \qquad q(s,\beta_i) = \int_{s_0}^{s} e^{-2\beta_i(s-u)}\sigma_i^2(u)\, du,$$

which completes the proof. □

Proof of Corollary 1 We first establish the limiting mean and covariance of $\boldsymbol{\mu}_i(s)$ as $s \to \infty$. We then use these limits to prove the convergence in distribution of $\boldsymbol{\mu}_i(s)$.

Step 1. By equation 19a,

$$\lim_{s\to\infty} \mathbf{m}_i(s) = \mathbf{m}_i, \tag{20}$$

since $e^{-\beta_i(s-s_0)} \to 0$ as $s \to \infty$.

Step 2. By equation 19b,

$$\begin{aligned}\lim_{s\to\infty} \Sigma_i(s) &= \lim_{s\to\infty} e^{-2\beta_i(s-s_0)}\Sigma_i + \mathrm{I}_d \lim_{s\to\infty} \int_{s_0}^{s} e^{-2\beta_i(s-u)}\sigma_i^2(u)\, du \\ &= \mathrm{I}_d \lim_{s\to\infty} q(s,\beta_i).\end{aligned} \tag{21}$$

The remaining limit can be written as

$$\begin{aligned}\lim_{s\to\infty} q(s,\beta_i) &= \lim_{s\to\infty} \frac{\int_{s_0}^{s} e^{2\beta_i u}\sigma_i^2(u)\, du}{e^{2\beta_i s}} \\ &= \lim_{s\to\infty} \frac{q_1(s,\beta_i)}{q_2(s,\beta_i)}.\end{aligned} \tag{22}$$

Clearly,

$$\lim_{s\to\infty} q_2(s,\beta_i) = \infty. \tag{23}$$

Since $\sigma_i^2(s) \to \sigma_i^2$, for any $\epsilon > 0$, there exists $u' > s_0$ such that for all $u \in [u', s]$,

$$\big|\sigma_i^2(u) - \sigma_i^2\big| < \epsilon, \tag{24}$$

which implies

$$\sigma_i^2 - \epsilon < \sigma_i^2(u) < \sigma_i^2 + \epsilon. \tag{25}$$

Therefore,

$$\begin{aligned}
\lim_{s\to\infty} q_1(s,\beta_i) &= \int_{s_0}^{u'} e^{2\beta_i u}\sigma_i^2(u)\,du + \lim_{s\to\infty}\int_{u'}^{s} e^{2\beta_i u}\sigma_i^2(u)\,du && \text{(integral decomposition)}\\
&\geq \lim_{s\to\infty}\int_{u'}^{s} e^{2\beta_i u}\sigma_i^2(u)\,du && \text{(discarding a nonnegative term)}\\
&> \lim_{s\to\infty}\frac{\sigma_i^2-\epsilon}{2\beta_i}\left(e^{2\beta_i s} - e^{2\beta_i u'}\right) && \text{(by equation 25)}\\
&> \lim_{s\to\infty}\frac{\sigma_i^2}{4\beta_i}\left(e^{2\beta_i s} - e^{2\beta_i u'}\right) && \text{(choosing } \epsilon = \sigma_i^2/2)\\
&= \infty. && (\sigma_i > 0)
\end{aligned} \tag{26}$$

Combining equation 23 and equation 26, L'Hôpital's rule yields

$$\begin{aligned}
\lim_{s\to\infty} q(s,\beta_i) &= \lim_{s\to\infty}\frac{\frac{d}{ds}\,q_1(s,\beta_i)}{\frac{d}{ds}\,q_2(s,\beta_i)} && \text{(L'Hôpital's rule)}\\
&= \lim_{s\to\infty}\frac{e^{2\beta_i s}\sigma_i^2(s)}{2\beta_i e^{2\beta_i s}} && \text{(Leibniz integral rule)}\\
&= \frac{\sigma_i^2}{2\beta_i}.
\end{aligned} \tag{27}$$

Substituting equation 27 into equation 21 gives

$$\lim_{s\to\infty}\Sigma_i(s) = \frac{\sigma_i^2}{2\beta_i}\mathrm{I}_d. \tag{28}$$

Step 3. Applying equation 20 and equation 28, for any $\mathbf{t}\in\mathbb{R}^d$, the characteristic function of $\boldsymbol{\mu}_i(s)$ satisfies

$$\begin{aligned}
\mathbb{E}\left(e^{\mathrm{i}\mathbf{t}^\top\boldsymbol{\mu}_i(s)}\right) &= \exp\left(\mathrm{i}\,\mathbf{t}^\top\mathbf{m}_i(s) - \frac{1}{2}\mathbf{t}^\top\Sigma_i(s)\mathbf{t}\right)\\
&\to \exp\left(\mathrm{i}\,\mathbf{t}^\top\mathbf{m}_i - \frac{1}{2}\mathbf{t}^\top\frac{\sigma_i^2\mathrm{I}_d}{2\beta_i}\mathbf{t}\right),
\end{aligned} \tag{29}$$

as $s\to\infty$, where $\mathrm{i}=\sqrt{-1}$ denotes the imaginary unit.

The limiting expression in equation 29 is the characteristic function of the Gaussian distribution with mean $\mathbf{m}_i$ and covariance $\frac{\sigma_i^2}{2\beta_i}\mathrm{I}_d$. Therefore, by Lévy's continuity theorem for characteristic functions (Van der Vaart, 2000, Theorem 2.13),

$$\boldsymbol{\mu}_i(s) \xrightarrow{d} \mathrm{N}\left(\mathbf{m}_i, \frac{\sigma_i^2}{2\beta_i}\mathrm{I}_d\right), \tag{30}$$

which completes the proof. □

Proof of Lemma 2 Recall that the finite mixture model in equation 5 is given by

$$p_N(\mathbf{x}(s)\mid G) = \sum_{j=1}^{N}\omega_j\,\mathrm{N}(\mathbf{x}(s);\boldsymbol{\mu}_j(s),\Sigma_j)\,, \tag{31a}$$

$$(\omega_1,\ldots,\omega_N)\sim\mathrm{Dir}(\alpha/N,\ldots,\alpha/N)\,. \tag{31b}$$

Also, for n i.i.d. stochastic processes $\{\mathbf{X}_i(s) : s \in \mathcal{T}\}_{i=1}^n$, let $Z_i(s) \in [\![N]\!]$ denote the latent allocation variable assigning $\mathbf{X}_i(s)$ to mixture component j. Define

$$I_j(s) = \{i : Z_i(s) = j\}, \qquad n_j(s) = \mathrm{card}(I_j(s)),$$

and let

$$\mathbf{X}_{I_j}(s) = \{\mathbf{X}_i(s) : i \in I_j(s)\}.$$

To prove part (i), the PP distribution is

$$\begin{aligned}
p_N(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) &= \int p_N(\mathbf{x}_{n+1}(s) \mid G)\, p_N(G \mid \mathbf{x}_{1:n}(s), z_{1:n}(s))\, dG \\
&= \iiint \sum_{j=1}^{N} w_j\, \mathrm{N}(\mathbf{x}_{n+1}(s); \boldsymbol{\mu}_j(s), \Sigma_j) \\
&\qquad \times p(w_{1:N}, \boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N} \mid \mathbf{x}_{1:n}(s), z_{1:n}(s))\; dw_{1:N}\, d\boldsymbol{\mu}_{1:N}(s)\, d\Sigma_{1:N} \\
&= \sum_{j=1}^{N} \Bigg[\iiint w_j\, \mathrm{N}(\mathbf{x}_{n+1}(s); \boldsymbol{\mu}_j(s), \Sigma_j) \\
&\qquad \times p(w_{1:N}, \boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N} \mid \mathbf{x}_{1:n}(s), z_{1:n}(s))\; dw_{1:N}\, d\boldsymbol{\mu}_{1:N}(s)\, d\Sigma_{1:N} \Bigg].
\end{aligned} \tag{32}$$

On the other hand,

$$\begin{aligned}
p(w_{1:N}, \boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N} \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) &= p(w_{1:N} \mid \boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N}, \mathbf{x}_{1:n}(s), z_{1:n}(s)) \\
&\qquad \times p(\boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N} \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) && \text{(33a)} \\
&= p(w_{1:N} \mid z_{1:n}(s)) \\
&\qquad \times p(\boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N} \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) && \text{(33b)} \\
&= p(w_{1:N} \mid z_{1:n}(s)) \\
&\qquad \times \prod_{k=1}^{N} p(\boldsymbol{\mu}_k(s), \Sigma_k \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) && \text{(33c)} \\
&= p(w_{1:N} \mid z_{1:n}(s)) \\
&\qquad \times \prod_{k=1}^{N} p(\boldsymbol{\mu}_k(s), \Sigma_k \mid \mathbf{x}_{I_k}(s)) && \text{(33d)} \\
&= p(w_{1:N} \mid z_{1:n}(s)) \\
&\qquad \times \prod_{k=1}^{N} \mathring{g}(\boldsymbol{\mu}_k(s), \Sigma_k \mid \mathbf{x}_{I_k}(s))\,, && \text{(33e)}
\end{aligned}$$

where equation 33b follows from the conditional independence

$$w_{1:N} \perp\!\!\!\perp (\boldsymbol{\mu}_{1:N}(s), \Sigma_{1:N}) \mid z_{1:n}(s).$$

The factorization in equation 33c follows from the conditional independence of the component-specific parameters given $(\mathbf{x}_{1:n}(s), z_{1:n}(s))$. Furthermore, equation 33d holds because the posterior distribution of the parameters associated with the k-th mixture component depends only on the observations assigned to that component. Finally, equation 33e merely introduces the notation.

Substituting equation 33 into equation 32 yields

$$
\begin{aligned}
p_N(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) = \sum_{j=1}^{N} \Bigg[& \int w_j \, p(w_{1:N} \mid z_{1:n}(s)) \; dw_{1:N} \\
& \times \iint \mathrm{N}(\mathbf{x}_{n+1}(s); \boldsymbol{\mu}_j(s), \Sigma_j) \, \mathring{g}\big(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\big) \; d\boldsymbol{\mu}_j(s) \, d\Sigma_j \\
& \times \prod_{k \neq j}^{N} \iint \mathring{g}(\boldsymbol{\mu}_k(s), \Sigma_k \mid \mathbf{x}_{I_k}(s)) \; d\boldsymbol{\mu}_k(s) \, d\Sigma_k \Bigg] \\
= \sum_{j=1}^{N} & \mathbb{E}(w_j \mid z_{1:n}(s)) \\
& \times \iint \mathrm{N}(\mathbf{x}_{n+1}(s); \boldsymbol{\mu}_j(s), \Sigma_j) \, \mathring{g}\big(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\big) \; d\boldsymbol{\mu}_j(s) \, d\Sigma_j \\
= \sum_{j=1}^{N} & \mathbb{E}(w_j \mid z_{1:n}(s)) \, \dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)), \qquad (34)
\end{aligned}
$$

where

$$
\prod_{k \neq j}^{N} \iint \mathring{g}(\boldsymbol{\mu}_k(s), \Sigma_k \mid \mathbf{x}_{I_k}(s)) \; d\boldsymbol{\mu}_k(s) \, d\Sigma_k = 1,
$$

since the integral of a probability density over its support is equal to one. This completes the proof. □

The next remark characterizes the asymptotic behavior of the PP equation 34, showing that it converges to the Chinese restaurant process (CRP) predictive distribution as $N \to \infty$ (Jordan, 2005), thereby preserving the classical BNP clustering structure.

Remark 2 *Under the assumptions of Lemma 2, let*

$$
O^+ = \{j : n_j(s) > 0\}, \qquad N^+ = \mathrm{card}(O^+).
$$

If $N \to \infty$ while $N^+/N \to 0$, then

$$
p_N(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) \longrightarrow \sum_{j \in O^+} \frac{n_j(s)}{n+\alpha} \, \dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)) + \frac{\alpha}{n+\alpha} \, \dot{g}(\mathbf{x}_{n+1}(s)), \quad (35)
$$

which reduces to the Chinese restaurant process (CRP) PP distribution, where

$$
\dot{g}(\mathbf{x}_{n+1}(s)) = \iint \mathrm{N}(\mathbf{x}_{n+1}(s) \mid \boldsymbol{\mu}(s), \Sigma) \, \mathring{g}(\boldsymbol{\mu}(s), \Sigma) \; d\boldsymbol{\mu}(s) \, d\Sigma
$$

is the prior predictive density for a new cluster.

Proof of Remark 3 We begin by deriving the posterior distribution of the mixture weights, which is used to compute $\mathbb{E}(w_j \mid z_{1:n}(s))$ in equation 34. Since

$$
p(z_{1:n}(s) \mid w_{1:N}) = \prod_{j=1}^{N} w_j^{n_j(s)}, \qquad p(w_{1:N}) \propto \prod_{j=1}^{N} w_j^{\alpha/N-1},
$$

Bayes' theorem yields

$$
\begin{aligned}
p(w_{1:N} \mid z_{1:n}(s)) & \propto p(z_{1:n}(s) \mid w_{1:N}) p(w_{1:N}) \\
& \propto \prod_{j=1}^{N} w_j^{n_j(s)} \prod_{j=1}^{N} w_j^{\alpha/N-1} = \prod_{j=1}^{N} w_j^{\alpha/N+n_j(s)-1}. \qquad (36)
\end{aligned}
$$

Hence,

$$
w_{1:N} \mid z_{1:n}(s) \sim \mathrm{Dir}\Big(\frac{\alpha}{N} + n_1(s), \ldots, \frac{\alpha}{N} + n_N(s)\Big), \qquad (37)
$$

and therefore

$$\mathbb{E}(w_j \mid z_{1:n}(s)) = \frac{\alpha/N + n_j(s)}{\alpha + n}. \tag{38}$$

Substituting equation 38 into equation 34 yields

$$\begin{aligned} p_N(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{1:n}(s), z_{1:n}(s)) &= \sum_{j \in O^+} \frac{\alpha/N + n_j(s)}{n + \alpha} \dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)) \\ &\quad + \sum_{j \notin O^+} \frac{\alpha/N}{n + \alpha} \dot{g}(\mathbf{x}_{n+1}(s)) \\ &= \sum_{j \in O^+} \frac{\alpha/N + n_j(s)}{n + \alpha} \dot{g}_j(\mathbf{x}_{n+1}(s) \mid \mathbf{x}_{I_j}(s)) \\ &\quad + \frac{\alpha(N - N^+)/N}{n + \alpha} \dot{g}(\mathbf{x}_{n+1}(s)). \end{aligned} \tag{39}$$

The terms indexed by $j \notin O^+$ correspond to assigning $\mathbf{x}_{n+1}(s)$ to a previously unoccupied component, i.e., creating a new cluster. Since these components have received no observations, their PP distribution equals the prior predictive density $\dot{g}(\mathbf{x}_{n+1}(s))$. Taking the limit $N \to \infty$ while $N^+/N \to 0$ in equation 39 completes the proof.

□

Proof of Proposition 1 For a fixed component j, define

$$\bar{\mathbf{x}}_j(s) = \frac{1}{n_j(s)} \sum_{i \in I_j(s)} \mathbf{x}_i(s), \qquad \mathrm{R}_j(s) = \sum_{i \in I_j(s)} \big(\mathbf{x}_i(s) - \bar{\mathbf{x}}_j(s)\big)\big(\mathbf{x}_i(s) - \bar{\mathbf{x}}_j(s)\big)^\top.$$

Up to a normalizing constant, the joint posterior density of $(\boldsymbol{\mu}_j(s), \Sigma_j)$ is

$$\begin{aligned} \mathring{g}\big(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\big) &\propto p\big(\mathbf{x}_{I_j}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j\big)\, \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j) \\ &= \prod_{i \in I_j(s)} p(\mathbf{x}_i(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j) && \text{(40a)} \\ &\quad \times \mathring{g}'(\boldsymbol{\mu}_j(s) \mid \Sigma_j) && \text{(40b)} \\ &\quad \times \mathring{g}''(\Sigma_j)\,. && \text{(40c)} \end{aligned}$$

Under the dynamic model of Section 4.1, the densities in equation 40a, equation 40b, and equation 40c correspond respectively to the distributions

$$\begin{aligned} \mathbf{X}_{I_j}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j &\sim \mathrm{N}(\boldsymbol{\mu}_j(s), \Sigma_j)\,, && \text{(41a)} \\ \boldsymbol{\mu}_j(s) \mid \Sigma_j &\sim \mathrm{N}(\mathbf{m}_j(s), \Sigma_j(s))\,, && \text{(41b)} \\ \Sigma_j &\sim \mathrm{IW}(\Lambda_0, \nu_0)\,. && \text{(41c)} \end{aligned}$$

Under the assumptions in equation 41, the posterior density in equation 40 becomes

$$\begin{aligned} \mathring{g}\left(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\right) &\propto |\Sigma_j|^{-\frac{n_j(s)}{2}} \exp\left\{-\frac{1}{2} \sum_{i \in I_j(s)} \big(\mathbf{x}_i(s) - \boldsymbol{\mu}_j(s)\big)^\top \Sigma_j^{-1} \big(\mathbf{x}_i(s) - \boldsymbol{\mu}_j(s)\big)\right\} \\ &\quad \times |\Sigma_j(s)|^{-\frac{1}{2}} \exp\left\{-\frac{1}{2}\big(\boldsymbol{\mu}_j(s) - \mathbf{m}_j(s)\big)^\top \Sigma_j^{-1}(s) \big(\boldsymbol{\mu}_j(s) - \mathbf{m}_j(s)\big)\right\} \\ &\quad \times |\Sigma_j|^{-\frac{\nu_0 + d + 1}{2}} \exp\left\{-\frac{1}{2} \operatorname{tr}\big(\Lambda_0 \Sigma_j^{-1}\big)\right\}. \end{aligned} \tag{42}$$

Using Murphy (2012, Eq. 4.195),

$$\begin{aligned} \sum_{i \in I_j(s)} \big(\mathbf{x}_i(s) - \boldsymbol{\mu}_j(s)\big)^\top \Sigma_j^{-1} \big(\mathbf{x}_i(s) - \boldsymbol{\mu}_j(s)\big) &= \operatorname{tr}\big(\mathrm{R}_j(s) \Sigma_j^{-1}\big) \\ &\quad + n_j(s) \big(\boldsymbol{\mu}_j(s) - \bar{\mathbf{x}}_j(s)\big)^\top \Sigma_j^{-1} \big(\boldsymbol{\mu}_j(s) - \bar{\mathbf{x}}_j(s)\big), \end{aligned} \tag{43}$$

and substituting this identity into equation 42 gives

$$\begin{aligned}\mathring{g}\left(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\right) \propto \exp\Bigg\{ & -\frac{1}{2} n_j(s) \big(\boldsymbol{\mu}_j(s) - \bar{\mathbf{x}}_j(s)\big)^\top \Sigma_j^{-1} \big(\boldsymbol{\mu}_j(s) - \bar{\mathbf{x}}_j(s)\big) \\ & - \frac{1}{2} \big(\boldsymbol{\mu}_j(s) - \mathbf{m}_j(s)\big)^\top \Sigma_j^{-1}(s) \big(\boldsymbol{\mu}_j(s) - \mathbf{m}_j(s)\big) \Bigg\} \\ & \times |\Sigma_j|^{-\frac{\nu_0 + n_j(s) + d + 1}{2}} |\Sigma_j(s)|^{-\frac{1}{2}} \\ & \times \exp\left\{ -\frac{1}{2} \operatorname{tr}\big[\big(\Lambda_0 + \mathrm{R}_j(s)\big)\Sigma_j^{-1}\big] \right\}. \end{aligned} \tag{44}$$

Collecting the terms involving $\boldsymbol{\mu}_j(s)$ in equation 44 and completing the square yields

$$\begin{aligned}\mathring{g}\left(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\right) & \propto \exp\left\{ -\frac{1}{2} \big(\boldsymbol{\mu}_j(s) - \mathbf{m}_{j,n_j}(s)\big)^\top \Sigma_{j,n_j}^{-1}(s) \big(\boldsymbol{\mu}_j(s) - \mathbf{m}_{j,n_j}(s)\big) \right\} \\ & \times \left[|\Sigma_j|^{-\frac{\nu_0 + n_j(s) + d + 1}{2}} |\Sigma_j(s)|^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2} \operatorname{tr}\big[\big(\Lambda_0 + \mathrm{R}_j(s)\big)\Sigma_j^{-1}\big] \right\} \right. \\ & \qquad \left. \times \exp\left\{ -\frac{1}{2} \mathrm{M}_j(s) \right\} \right] \\ & \propto \mathring{g}'\left(\boldsymbol{\mu}_j(s) \mid \Sigma_j, \mathbf{x}_{I_j}(s)\right) \mathring{g}''\left(\Sigma_j \mid \mathbf{x}_{I_j}(s)\right), \end{aligned} \tag{45}$$

where

$$\mathring{g}'\left(\boldsymbol{\mu}_j(s) \mid \Sigma_j, \mathbf{x}_{I_j}(s)\right) = \mathrm{N}\big(\boldsymbol{\mu}_j(s); \mathbf{m}_{j,n_j}(s), \Sigma_{j,n_j}(s)\big), \tag{46a}$$

$$\begin{aligned}\mathring{g}''\left(\Sigma_j \mid \mathbf{x}_{I_j}(s)\right) = |\Sigma_j|^{-\frac{\nu_0 + n_j(s) + d + 1}{2}} |\Sigma_j(s)|^{-\frac{1}{2}} \exp\left\{ -\frac{1}{2} \operatorname{tr}\big[\big(\Lambda_0 + \mathrm{R}_j(s)\big)\Sigma_j^{-1}\big] \right\} \\ \times \exp\left\{ -\frac{1}{2} \mathrm{M}_j(s) \right\}. \end{aligned} \tag{46b}$$

with

$$\mathbf{m}_{j,n_j}(s) = \Sigma_{j,n_j}(s) \Big(\Sigma_j^{-1}(s) \mathbf{m}_j(s) + n_j(s) \Sigma_j^{-1} \bar{\mathbf{x}}_j(s) \Big), \tag{47a}$$

$$\Sigma_{j,n_j}(s) = \Big(\Sigma_j^{-1}(s) + n_j(s) \Sigma_j^{-1} \Big)^{-1}, \tag{47b}$$

$$\mathrm{M}_j(s) = \mathbf{m}_j^\top(s) \Sigma_j^{-1}(s) \mathbf{m}_j(s) + n_j(s)\, \bar{\mathbf{x}}_j^\top(s) \Sigma_j^{-1} \bar{\mathbf{x}}_j(s) - \mathbf{m}_{j,n_j}^\top(s) \Sigma_{j,n_j}^{-1}(s) \mathbf{m}_{j,n_j}(s). \tag{47c}$$

Therefore, for $s > 0$, the posterior distribution

$$\mathring{g}\left(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s)\right)$$

admits a Gaussian conditional distribution for

$$\boldsymbol{\mu}_j(s) \mid \Sigma_j, \mathbf{x}_{I_j}(s),$$

whereas the marginal posterior of Σ_j is no longer inverse-Wishart.

However, at $s_0 = 0$,

$$c(s_0, \beta_j) = 1, \qquad q(s_0, \beta_j) = 0,$$

so that

$$\Sigma_j(s_0) = c(s_0, \beta_j) \Sigma_j + q(s_0, \beta_j) \mathrm{I}_d = \Sigma_j, \tag{48}$$

and

$$\Sigma_{j,n_j}^{-1}(s_0) = (n_j(s) + 1) \Sigma_j^{-1}.$$

Moreover, equation 47c simplifies to

$$\begin{aligned}
\mathrm{M}_j(s_0) &= \mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\mathbf{m}_j(s_0) + n_j(s)\,\bar{\mathbf{x}}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0) \\
&\quad - (n_j(s)+1)\left(\frac{\mathbf{m}_j(s_0)+n_j(s)\bar{\mathbf{x}}_j(s_0)}{n_j(s)+1}\right)^\top \Sigma_j^{-1}\left(\frac{\mathbf{m}_j(s_0)+n_j(s)\bar{\mathbf{x}}_j(s_0)}{n_j(s)+1}\right) \\
&= \mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\mathbf{m}_j(s_0) + n_j(s)\,\bar{\mathbf{x}}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0) \\
&\quad - \frac{1}{n_j(s)+1}\Big[\mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\mathbf{m}_j(s_0) + 2n_j(s)\,\mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0) + n_j(s)^2\,\bar{\mathbf{x}}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0)\Big] \\
&= \frac{n_j(s)}{n_j(s)+1}\Big[\mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\mathbf{m}_j(s_0) - 2\mathbf{m}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0) + \bar{\mathbf{x}}_j^\top(s_0)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s_0)\Big] \\
&= \frac{n_j(s)}{n_j(s)+1}\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big)^\top\Sigma_j^{-1}\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big) \\
&= \frac{n_j(s)}{n_j(s)+1}\,\mathrm{tr}\Big[\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big)\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big)^\top\Sigma_j^{-1}\Big]. \qquad (49)
\end{aligned}$$

Combining equation 48, equation 49, and equation 46b yields

$$\begin{aligned}
\mathring{g}''\left(\Sigma_j \mid \mathbf{x}_{I_j}(s_0)\right) &= |\Sigma_j|^{-\frac{(\nu_0+n_j(s)+1)+d+1}{2}} \\
&\quad \times \exp\left\{-\frac{1}{2}\,\mathrm{tr}\left[\left(\Lambda_0 + \mathrm{R}_j(s_0) + \frac{n_j(s)}{n_j(s)+1}\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big)\big(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0)\big)^\top\right)\Sigma_j^{-1}\right]\right\}. \qquad (50)
\end{aligned}$$

Hence,

$$\mathring{g}''\left(\Sigma_j \mid \mathbf{x}_{I_j}(s_0)\right)$$

corresponds to

$$\mathrm{IW}\left(\nu_0 + n_j(s) + 1,\, \Lambda_0 + \mathrm{R}_j(s_0) + \frac{n_j(s)}{n_j(s)+1}(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0))(\bar{\mathbf{x}}_j(s_0)-\mathbf{m}_j(s_0))^\top\right).$$

□

Proof of Theorem 1 Recall that $\mathbf{m}_j(s)$ and $\mathbf{m}_{j,n_j}(s)$ in equation 4 and equation 8 are given by

$$\begin{aligned}
\mathbf{m}_j(s) &= \mathbf{m}_j + e^{-s\beta_j}(\mathbf{m}_j(s_0)-\mathbf{m}_j), \\
\mathbf{m}_{j,n_j}(s) &= \Sigma_{j,n_j}(s)\Big(\Sigma_j^{-1}(s)\mathbf{m}_j(s) + n_j(s)\Sigma_j^{-1}\bar{\mathbf{x}}_j(s)\Big).
\end{aligned}$$

Proof of equation 9 in Theorem 1(i): Using the expressions for $\mathbf{m}_j(s)$ and $\mathbf{m}_{j,n_j}(s)$ above, we obtain

$$\|\mathbf{m}_{j,n_j}(s)-\mathbf{m}_j(s)\|_2 = \left\|\left(\mathbf{I}_d - \boldsymbol{\Sigma}_{j,n_j}(s)\boldsymbol{\Sigma}_j^{-1}(s)\right)\mathbf{m}_j(s) - \boldsymbol{\Sigma}_{j,n_j}(s)\,n_j(s)\,\boldsymbol{\Sigma}_j^{-1}\,\bar{\mathbf{x}}_j(s)\right\|_2. \qquad (51)$$

For any fixed $s \in \mathcal{T}$ and $\beta_j > 0$, define

$$\Psi(s,\beta_j) := \Sigma_{j,n_j}(s)\Sigma_j^{-1}(s), \qquad (52a)$$

$$\mathrm{B} := n_j(s)\Sigma_j^{-1}, \qquad (52b)$$

$$\Sigma_{j,n_j}(s) := \left(\Sigma_j^{-1}(s) + \mathrm{B}\right)^{-1}. \qquad (52c)$$

We obtain

$$\begin{aligned}
\Psi(s,\beta_j) &= \left(\Sigma_j^{-1}(s) + \mathrm{B}\right)^{-1}\Sigma_j^{-1}(s), \\
&= \left(\mathrm{I}_d + \Sigma_j(s)\,\mathrm{B}\right)^{-1}, \qquad (53)
\end{aligned}$$

and

$$\begin{aligned} \mathrm{I}_d - \Psi(s, \beta_j) &= \left(\Sigma_j^{-1}(s) + \mathrm{B}\right)^{-1} \mathrm{B}, \\ &= \Sigma_{j,n_j}(s)\, \mathrm{B}, \\ &= \Sigma_{j,n_j}(s)\, n_j(s)\, \Sigma_j^{-1}, \end{aligned} \tag{54}$$

where equation 54 follows from the identity

$$\left(\Sigma_j^{-1}(s) + \mathrm{B}\right)^{-1} \left(\Sigma_j^{-1}(s) + \mathrm{B}\right) = \mathrm{I}_d.$$

Applying equation 52a and equation 54 to equation 51, we have

$$\begin{aligned} \|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 &= \|(\mathrm{I}_d - \Psi(s, \beta_j))\,(\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s))\|_2 \\ &\leq \|\mathrm{I}_d - \Psi(s, \beta_j)\|_{\mathrm{op}}\, \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2\,, \end{aligned} \tag{55}$$

where

$$\begin{aligned} \|\mathrm{I}_d - \Psi(s, \beta_j)\|_{\mathrm{op}} &= \max_{\|\mathbf{v}\|_2 = 1} \|(\mathrm{I}_d - \Psi(s, \beta_j))\,\mathbf{v}\|_2 \\ &= \sqrt{\max_{i \in [\![d]\!]} \lambda_i\Big(\left(\mathrm{I}_d - \Psi(s, \beta_j)\right)^\top \left(\mathrm{I}_d - \Psi(s, \beta_j)\right)\Big)}, \end{aligned} \tag{56}$$

with $\|\cdot\|_{\mathrm{op}}$ denoting the induced (operator) 2-norm and $\lambda_i(\cdot)$ denoting the i-th nonnegative eigenvalue of its matrix argument for $i \in [\![d]\!]$ (Meyer, 2000, Eq. 5.2.7).

Furthermore, since Σ_j is symmetric and positive semidefinite, by Horn & Johnson (1990, Corollary 2.5.14) there exists an orthogonal matrix Q such that

$$\Sigma_j = \mathrm{Q}\,\Lambda\,\mathrm{Q}^\top, \tag{57}$$

where Λ is a diagonal matrix whose entries are the nonnegative eigenvalues of Σ_j.

Substituting equation 57 into equation 52c and equation 52b, both $\Sigma_j(s)$ and B can be expressed in the same eigenbasis:

$$\Sigma_j(s) = \mathrm{Q}\big(c(s, \beta_j)\Lambda + q(s, \beta_j)\mathrm{I}_d\big)\mathrm{Q}^\top, \tag{58a}$$

$$\mathrm{B} = \mathrm{Q}\big(n_j(s)\Lambda^{-1}\big)\mathrm{Q}^\top. \tag{58b}$$

Therefore,

$$\begin{aligned} \Psi^{-1}(s, a) &= \mathrm{I}_d + \Sigma_j(s)\mathrm{B} && \text{(by equation 53)} \\ &= \mathrm{I}_d + \mathrm{Q}\big(c(s, \beta_j)\Lambda + q(s, \beta_j)\mathrm{I}_d\big)\mathrm{Q}^\top \mathrm{Q}\big(n_j(s)\Lambda^{-1}\big)\mathrm{Q}^\top && \text{(by equation 58)} \\ &= \mathrm{Q}\mathrm{Q}^\top + \mathrm{Q}\big(c(s, \beta_j)\Lambda + q(s, \beta_j)\mathrm{I}_d\big)\big(n_j(s)\Lambda^{-1}\big)\mathrm{Q}^\top && (\mathrm{Q}^\top\mathrm{Q} = \mathrm{I}_d) \\ &= \mathrm{Q}\left(\mathrm{I}_d + (c(s, \beta_j)\Lambda + q(s, \beta_j)\mathrm{I}_d) n_j(s)\Lambda^{-1}\right)\mathrm{Q}^\top && \text{(algebra)} \\ &= \mathrm{Q}\,\mathrm{diag}\Big(1 + n_j(s)\left(c(s, \beta_j) + \tfrac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\right)\Big)\,\mathrm{Q}^\top && \text{(notation)}, \end{aligned} \tag{59}$$

where $\lambda_i(\Sigma_j)$ denotes the i-th nonnegative eigenvalue of Σ_j for $i \in [\![d]\!]$.

By the orthogonality property $\mathrm{Q}^\top = \mathrm{Q}^{-1}$, taking the inverse of equation 59 gives

$$\Psi(s, a) = \mathrm{Q}\,\mathrm{diag}\left(\frac{1}{1 + n_j(s)\left(c(s, \beta_j) + \frac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\right)}\right)\mathrm{Q}^\top.$$

Applying Horn & Johnson (1990, Corollary 2.5.14) again to this decomposition, it follows that $\Psi(s, a)$ is symmetric, i.e., $\Psi(s, a) = \Psi(s, a)^\top$, with eigenvalues

$$\lambda_i(\Psi(s, \beta_j)) = \frac{1}{1 + n_j(s)\left(c(s, \beta_j) + \frac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\right)} \in [0, 1]. \tag{60}$$

Applying the above properties to equation 56 implies

$$\begin{aligned}
\|\mathrm{I}_d - \Psi(s,\beta_j)\|_{\mathrm{op}} &= \sqrt{\max_{i\in[\![d]\!]} \lambda_i((\mathrm{I}_d - \Psi(s,\beta_j))^2)} && (\mathrm{I}_d - \Psi(s,\beta_j) \text{ is symmetric}) \\
&= \sqrt{\max_{i\in[\![d]\!]} \lambda_i^2(\mathrm{I}_d - \Psi(s,\beta_j))} && \text{(eigenvalue property)} \\
&= \sqrt{\max_{i\in[\![d]\!]} \left(1 - \lambda_i(\Psi(s,\beta_j))\right)^2} && \text{(eigenvalue property)} \\
&= \max_{i\in[\![d]\!]} \left(1 - \lambda_i(\Psi(s,\beta_j))\right) && (\lambda_i(\Psi(s,\beta_j)) \in [0,1]) \\
&= \max_{i\in[\![d]\!]} \left(1 - \frac{1}{1 + n_j(s)\Big(c(s,\beta_j) + \frac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\Big)}\right) && \text{(by equation 60)} \\
&= \max_{i\in[\![d]\!]} \frac{n_j(s)\Big(c(s,\beta_j) + \frac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\Big)}{1 + n_j(s)\Big(c(s,\beta_j) + \frac{q(s,\beta_j)}{\lambda_i(\Sigma_j)}\Big)} && \text{(algebra)} \\
&= \frac{n_j(s)\Big(c(s,\beta_j) + \frac{q(s,\beta_j)}{\varsigma_j^2}\Big)}{1 + n_j(s)\Big(c(s,\beta_j) + \frac{q(s,\beta_j)}{\varsigma_j^2}\Big)} && (\Sigma_j = \varsigma_j^2 \mathrm{I}_d). \\
&= \mathcal{U}_s(\beta_j) && \text{(notation).}
\end{aligned} \tag{61}$$

Applying equation 61 to equation 51 establishes equation 9 in Theorem 1(i):

$$\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \le \mathcal{U}_s(\beta_j)\|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2.$$

Proof of equation 10 in Theorem 1(i): We obtain

$$\begin{aligned}
\|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 &= \left\| \left(\boldsymbol{\Sigma}_{j,n_j}(s)\boldsymbol{\Sigma}_j^{-1}(s)\right) \mathbf{m}_j(s) - \left(\mathrm{I}_d - \boldsymbol{\Sigma}_{j,n_j}(s)\, n_j(s)\, \boldsymbol{\Sigma}_j^{-1}\right)\, \bar{\mathbf{x}}_j(s) \right\|_2 \\
&= \|\Psi(s,\beta_j)\,(\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s))\|_2 && \text{(by equation 52a \& equation 54)} \\
&\le \|\Psi(s,\beta_j)\|_{\mathrm{op}}\, \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 && \text{(operator norm definition)} \\
&= (1 - \mathcal{U}_s(\beta_j))\, \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 && \text{(similar to equation 61).}
\end{aligned} \tag{62}$$

This completes the proof of part (i).

Proof of Theorem 1(ii): Note that since $\Sigma_j(s)$ and Σ_j are positive semidefinite,

$$\Sigma_j^{-1}(s) + n_j(s)\Sigma_j^{-1}(s_0) \succeq \Sigma_j^{-1}(s), \qquad \Sigma_j^{-1}(s) + n_j(s)\Sigma_j^{-1}(s_0) \succeq n_j(s)\Sigma_j^{-1}(s_0),$$

where "$\succeq$" denotes the Loewner (semidefinite) order, i.e., for symmetric matrices A and B, $A \succeq B$ if and only if $A - B$ is positive semidefinite. Since matrix inversion reverses this order on the positive semidefinite cone, it follows that

$$\Sigma_{j,n_j}(s) = \left(\Sigma_j^{-1}(s) + n_j(s)\Sigma_j^{-1}(s_0)\right)^{-1} \preceq \Sigma_j(s), \qquad \Sigma_{j,n_j}(s) \preceq \tfrac{1}{n_j(s)}\Sigma_j.$$

By the definition of positive semidefiniteness, for all $\mathbf{v} \neq 0$,

$$\dot{\mathbf{v}}^\top \Sigma_{j,n_j}(s)\, \dot{\mathbf{v}} \le \dot{\mathbf{v}}^\top \Sigma_j(s)\, \dot{\mathbf{v}}, \qquad \dot{\mathbf{v}}^\top \Sigma_{j,n_j}(s)\, \dot{\mathbf{v}} \le \dot{\mathbf{v}}^\top \left(\tfrac{1}{n_j(s)}\Sigma_j\right) \dot{\mathbf{v}} \le \dot{\mathbf{v}}^\top \Sigma_j \dot{\mathbf{v}},$$

where $\dot{\mathbf{v}} = \mathbf{v}/\|\mathbf{v}\|_2$ denotes the normalized vector. Taking the maximum over all such $\dot{\mathbf{v}}$, Meyer (2000, Example 7.5.1) implies that

$$\begin{aligned}
\lambda_{\max}\big(\Sigma_{j,n_j}(s)\big) &= \max_{\dot{\mathbf{v}}} \dot{\mathbf{v}}^\top \Sigma_{j,n_j}(s)\, \dot{\mathbf{v}} \le \max_{\dot{\mathbf{v}}} \dot{\mathbf{v}}^\top \Sigma_j(s)\, \dot{\mathbf{v}} = \lambda_{\max}\big(\Sigma_j(s)\big), \\
\lambda_{\max}\big(\Sigma_{j,n_j}(s)\big) &= \max_{\dot{\mathbf{v}}} \dot{\mathbf{v}}^\top \Sigma_{j,n_j}(s)\, \dot{\mathbf{v}} \le \max_{\dot{\mathbf{v}}} \dot{\mathbf{v}}^\top \Sigma_j \dot{\mathbf{v}} = \lambda_{\max}\big(\Sigma_j\big),
\end{aligned}$$

which completes the proof. □

Proof of Corollary 2 Recall from Theorem 1(i) that

$$\|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 \leq \mathcal{U}_s(\beta_j)\|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2, \tag{63a}$$

$$\|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \leq \big(1 - \mathcal{U}_s(\beta_j)\big)\|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2. \tag{63b}$$

Here,

$$\mathcal{U}_s(\beta_j) = \frac{n_j(s)\left(c(s,\beta_j) + \dfrac{q(s,\beta_j)}{\varsigma_j^2}\right)}{1 + n_j(s)\left(c(s,\beta_j) + \dfrac{q(s,\beta_j)}{\varsigma_j^2}\right)}, \tag{64a}$$

$$c(s,\beta_j) = e^{-2s\beta_j}, \tag{64b}$$

$$q(s,\beta_j) = \frac{\int_0^s e^{2\beta_j u}\sigma_j^2(u)\,du}{e^{2\beta_j s}} =: \frac{q_1(s,\beta_j)}{q_2(s,\beta_j)}. \tag{64c}$$

To establish the result, we first study the asymptotic behavior of $\mathcal{U}_s(\beta_j)$ as $s \to \infty$. Since equation 64b immediately yields

$$\lim_{s\to\infty} c(s,\beta_j) = 0, \tag{65}$$

it remains to analyze the convergence of $q(s,\beta_j)$, which depends on the growth rate of $\sigma_j^2(s)$.

For part (i), assume that $\sigma_j^2(s) \to \infty$. Then, for any $M > 0$, there exists $u' > 0$ such that

$$\sigma_j^2(u) > M, \qquad u \in [u', s]. \tag{66}$$

Consequently,

$$\begin{aligned}
\lim_{s\to\infty} q_1(s,\beta_j) &= \int_0^{u'} e^{2\beta_j u}\sigma_j^2(u)\,du + \lim_{s\to\infty}\int_{u'}^{s} e^{2\beta_j u}\sigma_j^2(u)\,du \\
&\geq \lim_{s\to\infty}\int_{u'}^{s} e^{2\beta_j u}\sigma_j^2(u)\,du \\
&> \lim_{s\to\infty} M\int_{u'}^{s} e^{2\beta_j u}\,du \qquad \text{(by equation 66)} \\
&= \lim_{s\to\infty}\frac{M}{2\beta_j}\left(e^{2\beta_j s} - e^{2\beta_j u'}\right) \\
&= \infty.
\end{aligned} \tag{67}$$

Since

$$\lim_{s\to\infty} q_2(s,\beta_j) = \lim_{s\to\infty} e^{2\beta_j s} = \infty, \tag{68}$$

the limit of equation 64c is of the indeterminate form ∞/∞. Hence, by L'Hôpital's rule,

$$\begin{aligned}
\lim_{s\to\infty} q(s,\beta_j) &= \lim_{s\to\infty}\frac{e^{2\beta_j s}\sigma_j^2(s)}{2\beta_j e^{2\beta_j s}} \\
&= \lim_{s\to\infty}\frac{\sigma_j^2(s)}{2\beta_j} = \infty.
\end{aligned} \tag{69}$$

Substituting this result together with $\lim_{s\to\infty} c(s,\beta_j) = 0$ into equation 64a gives

$$\lim_{s\to\infty} \mathcal{U}_s(\beta_j) = 1. \tag{70}$$

Taking $s \to \infty$ in equation 63b yields

$$\begin{aligned}
0 &\leq \lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \\
&\leq \lim_{s\to\infty} \big(1 - \mathcal{U}_s(\beta_j)\big) \lim_{s\to\infty} \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 \\
&= 0 \times \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2 \qquad \text{(by equation 70 and the continuous mapping theorem)} \\
&= 0. \qquad (71)
\end{aligned}$$

Therefore,

$$\lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 = 0. \qquad (72)$$

Next, taking $s \to \infty$ in equation 63a gives

$$\begin{aligned}
\lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 &\leq \lim_{s\to\infty} \mathcal{U}_s(\beta_j) \lim_{s\to\infty} \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 \\
&= \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2. \qquad (73)
\end{aligned}$$

Moreover,

$$\begin{aligned}
\|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2 &= \lim_{s\to\infty} \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 \qquad \text{(continuous mapping theorem)} \\
&\leq \lim_{s\to\infty} \|\mathbf{m}_j(s) - \mathbf{m}_{j,n_j}(s)\|_2 + \lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \\
&\qquad \text{(triangle inequality)} \\
&= \lim_{s\to\infty} \|\mathbf{m}_j(s) - \mathbf{m}_{j,n_j}(s)\|_2 \qquad \text{(by equation 72).} \\
&\qquad (74)
\end{aligned}$$

Combining equation 73 and equation 74 yields

$$\lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 = \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2. \qquad (75)$$

The result in part (i) now follows directly from equation 72 and equation 75.

To prove part (ii), assume that $\sigma_j^2(s) \to 0$. Then, for any $\epsilon > 0$, there exists $u' > 0$ such that

$$\sigma_j^2(u) < \epsilon, \qquad u \in [u', s]. \qquad (76)$$

Hence,

$$\begin{aligned}
\lim_{s\to\infty} q(s, \beta_j) &= \lim_{s\to\infty} \frac{\int_0^{u'} e^{2\beta_j u} \sigma_j^2(u)\, du}{e^{2\beta_j s}} + \lim_{s\to\infty} \frac{\int_{u'}^{s} e^{2\beta_j u} \sigma_j^2(u)\, du}{e^{2\beta_j s}} \\
&= 0 + \lim_{s\to\infty} \frac{\int_{u'}^{s} e^{2\beta_j u} \sigma_j^2(u)\, du}{e^{2\beta_j s}} \\
&< \lim_{s\to\infty} \frac{\epsilon \int_{u'}^{s} e^{2\beta_j u}\, du}{e^{2\beta_j s}} \qquad \text{(by equation 76)} \\
&= \lim_{s\to\infty} \frac{\epsilon \left(e^{2\beta_j s} - e^{2\beta_j u'}\right)}{2\beta_j e^{2\beta_j s}} \\
&= \frac{\epsilon}{2\beta_j}. \qquad (77)
\end{aligned}$$

Since $q(s, \beta_j) \geq 0$ and $\epsilon > 0$ is arbitrary, it follows that

$$\lim_{s\to\infty} q(s, \beta_j) = 0. \qquad (78)$$

Combining this result with

$$\lim_{s\to\infty} c(s, \beta_j) = 0, \qquad (79)$$

and using equation 64a, we obtain

$$\lim_{s\to\infty} \mathcal{U}_s(\beta_j) = 0. \tag{80}$$

Taking $s \to \infty$ in equation 63a yields

$$\lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \mathbf{m}_j(s)\|_2 = 0. \tag{81}$$

Moreover,

$$\begin{aligned}
\|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2 &= \lim_{s\to\infty} \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 && \text{(continuous mapping theorem)} \\
&\le \lim_{s\to\infty} \|\mathbf{m}_j(s) - \mathbf{m}_{j,n_j}(s)\|_2 + \lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 \\
& && \text{(triangle inequality)} \\
&= \lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 && \text{(by equation 81)} \\
&\le \lim_{s\to\infty} \big(1 - \mathcal{U}_s(\beta_j)\big) \lim_{s\to\infty} \|\mathbf{m}_j(s) - \bar{\mathbf{x}}_j(s)\|_2 && \text{(by equation 63b)} \\
&= \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2 && \text{(by equation 80 and continuous mapping theorem).}
\end{aligned} \tag{82}$$

Therefore,

$$\lim_{s\to\infty} \|\mathbf{m}_{j,n_j}(s) - \bar{\mathbf{x}}_j(s)\|_2 = \|\mathbf{m}_j - \bar{\mathbf{x}}_j\|_2. \tag{83}$$

Combining equation 81 and equation 83 completes the proof of part (ii).

To prove part (iii), suppose that $\sigma_j^2(u) \le M < \infty$ for all $u \in [0, s]$. Then,

$$\begin{aligned}
\lim_{\beta_j\to\infty} q(s, \beta_j) &= \lim_{\beta_j\to\infty} \frac{\int_0^s e^{2\beta_j u}\sigma_j^2(u)\, du}{e^{2\beta_j s}} \\
&\le \lim_{\beta_j\to\infty} \frac{M \int_0^s e^{2\beta_j u}\, du}{e^{2\beta_j s}} \\
&= \lim_{\beta_j\to\infty} \frac{M\left(e^{2\beta_j s} - 1\right)}{2\beta_j e^{2\beta_j s}} \\
&= \lim_{\beta_j\to\infty} \frac{M\left(1 - e^{-2\beta_j s}\right)}{2\beta_j} \\
&= 0.
\end{aligned} \tag{84}$$

Since $q(s, \beta_j) \ge 0$, equation 84 implies

$$\lim_{\beta_j\to\infty} q(s, \beta_j) = 0. \tag{85}$$

The remainder of the proof follows the same argument as in part (ii). In particular, replacing the limit $s \to \infty$ in equation 80–equation 83 by $\beta_j \to \infty$ yields the desired result. □

Proof of Theorem 2 Recall from equation 37 that, conditional on the latent allocation vectors $z_{1:n}(s)$ and $z'_{1:n}(s)$, with $z_i(s) = j$ and $z'_i(s) = j$ indicating that $\mathbf{x}_i(s)$ and $\mathbf{x}'_i(s)$ belong to mixture

component j, respectively,

$$\log \frac{p(w_{1:N} \mid z_{1:n}(s))}{p(w_{1:N} \mid z'_{1:n}(s))} = \log \frac{\frac{\Gamma(\alpha+n)}{\prod_{j=1}^{N}\Gamma(\alpha/N+n\pi_j(s))}\prod_{j=1}^{N} w_j^{\alpha/N+n\pi_j(s)-1}}{\frac{\Gamma(\alpha+n)}{\prod_{j=1}^{N}\Gamma(\alpha/N+n\pi'_j(s))}\prod_{j=1}^{N} w_j^{\alpha/N+n\pi'_j(s)-1}}$$
$$= \underbrace{\sum_{j=1}^{N}\Big[\log\Gamma\Big(\frac{\alpha}{N}+n\pi'_j(s)\Big) - \log\Gamma\Big(\frac{\alpha}{N}+n\pi_j(s)\Big)\Big]}_{\mathcal{I}_1}$$
$$+ \underbrace{n\sum_{j=1}^{N}(\pi_j(s)-\pi'_j(s))\log w_j}_{\mathcal{I}_2}. \tag{86}$$

We now bound the two terms in equation 86, starting with $\mathcal{I}_1$. For a fixed j, first suppose that $\pi_j(s) \le \pi'_j(s)$. By the mean value theorem, there exists

$$\xi_j \in \Big[\frac{\alpha}{N}+n\pi_j(s),\ \frac{\alpha}{N}+n\pi'_j(s)\Big] \tag{87}$$

such that

$$\frac{d}{d\xi_j}\log\Gamma(\xi_j) = \frac{\log\Gamma\big(\frac{\alpha}{N}+n\pi'_j(s)\big) - \log\Gamma\big(\frac{\alpha}{N}+n\pi_j(s)\big)}{n\big(\pi'_j(s)-\pi_j(s)\big)}. \tag{88}$$

Using

$$\psi(\xi_j) = \frac{d}{d\xi_j}\log\Gamma(\xi_j), \tag{89}$$

where ψ denotes the digamma function, equation 88 gives

$$\log\Gamma\Big(\frac{\alpha}{N}+n\pi'_j(s)\Big) - \log\Gamma\Big(\frac{\alpha}{N}+n\pi_j(s)\Big) = \psi(\xi_j)n\big(\pi'_j(s)-\pi_j(s)\big). \tag{90}$$

Summing equation 90 over j yields

$$\mathcal{I}_1 = \sum_{j=1}^{N}\Big[\log\Gamma\Big(\frac{\alpha}{N}+n\pi'_j(s)\Big) - \log\Gamma\Big(\frac{\alpha}{N}+n\pi_j(s)\Big)\Big]$$
$$= \sum_{j=1}^{N}\psi(\xi_j)n\big(\pi'_j(s)-\pi_j(s)\big). \tag{91}$$

Since $\sum_{j=1}^{N}\big(\pi'_j(s)-\pi_j(s)\big) = 0$, we may subtract the common term $\psi(\alpha/N)$ to obtain

$$\mathcal{I}_1 = \sum_{j=1}^{N} n\left[\psi(\xi_j) - \psi\left(\frac{\alpha}{N}\right)\right]\big(\pi'_j(s)-\pi_j(s)\big). \tag{92}$$

Taking absolute values and applying the triangle inequality gives

$$|\mathcal{I}_1| \le n\sum_{j=1}^{N}\left|\psi(\xi_j) - \psi\left(\frac{\alpha}{N}\right)\right|\left|\pi'_j(s)-\pi_j(s)\right|. \tag{93}$$

Since $0 \le \pi_j(s), \pi'_j(s) \le 1$, we have

$$\frac{\alpha}{N} \le \xi_j \le \frac{\alpha}{N}+n. \tag{94}$$

Moreover, since ψ is increasing on $(0,\infty)$,

$$0 \le \psi(\xi_j) - \psi\left(\frac{\alpha}{N}\right) \le \psi\left(\frac{\alpha}{N}+n\right) - \psi\left(\frac{\alpha}{N}\right). \tag{95}$$

Substituting equation 95 into equation 93 yields

$$\begin{aligned}|\mathcal{I}_1| &\le n\left[\psi\left(\frac{\alpha}{N}+n\right)-\psi\left(\frac{\alpha}{N}\right)\right]\sum_{j=1}^{N}\left|\pi'_j(s)-\pi_j(s)\right| \\ &= n\left[\psi\left(\frac{\alpha}{N}+n\right)-\psi\left(\frac{\alpha}{N}\right)\right]\left\|\boldsymbol{\pi}(\mathbf{X}_{1:n}(s))-\boldsymbol{\pi}(\mathbf{X}'_{1:n}(s))\right\|_1 \\ &\le 2\left[\psi\left(\frac{\alpha}{N}+n\right)-\psi\left(\frac{\alpha}{N}\right)\right]. \end{aligned} \tag{96}$$

By the mean value theorem, there exists $\zeta \in [\frac{\alpha}{N}, \frac{\alpha}{N}+n]$ such that

$$\psi\left(\frac{\alpha}{N}+n\right)-\psi\left(\frac{\alpha}{N}\right) = n\psi'(\zeta), \tag{97}$$

where ψ' denotes the trigamma function. Therefore,

$$|\mathcal{I}_1| \le 2n \max_{\zeta\in\left[\frac{\alpha}{N},\,\frac{\alpha}{N}+n\right]} \psi'(\zeta). \tag{98}$$

Since ψ' is strictly decreasing on $(0,\infty)$, its maximum over $\left[\frac{\alpha}{N}, \frac{\alpha}{N}+n\right]$ is attained at the left endpoint. Hence,

$$\mathcal{I}_1 \le |\mathcal{I}_1| \le 2n\psi'\left(\frac{\alpha}{N}\right). \tag{99}$$

We next bound $\mathcal{I}_2$ in equation 86. Let

$$d_j(s) := \pi_j(s) - \pi'_j(s), \qquad j \in [\![N]\!]. \tag{100}$$

Since $\boldsymbol{\pi}(\mathbf{X}_{1:n}(s)), \boldsymbol{\pi}(\mathbf{X}'_{1:n}(s)) \in \Omega^{\pi}_{\gamma_\alpha} \subset \Delta_N$,

$$\sum_{j=1}^{N} d_j(s) = 0. \tag{101}$$

Define

$$D^+ := \{j : d_j(s) \ge 0\}, \qquad D^- := \{j : d_j(s) < 0\}. \tag{102}$$

Then

$$\sum_{j\in D^+} d_j(s) = \sum_{j\in D^-}\big(-d_j(s)\big) = \frac{1}{2}\left\|\boldsymbol{\pi}(\mathbf{X}_{1:n}(s)) - \boldsymbol{\pi}(\mathbf{X}'_{1:n}(s))\right\|_1 \le \frac{1}{n}. \tag{103}$$

We first restrict attention to the event $\{w_{1:N} \in \Omega^{w}_{\gamma_\alpha}\}$. On this event, since $\sum_{j=1}^{N} w_j = 1$, we have, for every $j \in [\![N]\!]$,

$$\gamma_\alpha \le w_j \le 1-(N-1)\gamma_\alpha, \tag{104}$$

and therefore

$$\log\gamma_\alpha \le \log w_j \le \log\big(1-(N-1)\gamma_\alpha\big). \tag{105}$$

Hence,

$$\begin{aligned}
\mathcal{I}_2 &= n\sum_{j=1}^{N} d_j(s)\log w_j \\
&= n\sum_{j\in D^+} d_j(s)\log w_j + n\sum_{j\in D^-} d_j(s)\log w_j \\
&\le n\log\big(1-(N-1)\gamma_\alpha\big)\sum_{j\in D^+} d_j(s) + n\log\gamma_\alpha \sum_{j\in D^-} d_j(s) \\
&= n\log\big(1-(N-1)\gamma_\alpha\big)\sum_{j\in D^+} d_j(s) - n\log\gamma_\alpha \sum_{j\in D^-} \big(-d_j(s)\big) \\
&= n\big[\log\big(1-(N-1)\gamma_\alpha\big) - \log\gamma_\alpha\big]\sum_{j\in D^+} d_j(s) \\
&= n\log\left(\frac{1-(N-1)\gamma_\alpha}{\gamma_\alpha}\right)\sum_{j\in D^+} d_j(s) \\
&\le \log\left(\frac{1-(N-1)\gamma_\alpha}{\gamma_\alpha}\right), \qquad (106)
\end{aligned}$$

where the last inequality follows from equation 103.

Combining equation 86, equation 99, and equation 106 yields, on the event $\{w_{1:N} \in \Omega^w_{\gamma_\alpha}\}$,

$$\log\frac{p(w_{1:N} \mid z_{1:n}(s))}{p(w_{1:N} \mid z'_{1:n}(s))} \le 2n\psi'\left(\frac{\alpha}{N}\right) + \log\left(\frac{1-(N-1)\gamma_\alpha}{\gamma_\alpha}\right) = \epsilon(\alpha). \qquad (107)$$

Thus, the likelihood-ratio bound holds on $\Omega^w_{\gamma_\alpha}$.

By the definition of $\delta(\alpha)$, for every $\boldsymbol{\pi} \in \Omega^\pi_{\gamma_\alpha}$,

$$\Pr\big[\mathcal{M}(\boldsymbol{\pi}) \in (\Omega^w_{\gamma_\alpha})^c\big] \le \delta(\alpha). \qquad (108)$$

Therefore, for any measurable $A \subseteq \Delta_N$,

$$\begin{aligned}
\Pr[\mathcal{M}(\boldsymbol{\pi}) \in A] &\le \Pr[\mathcal{M}(\boldsymbol{\pi}) \in A \cap \Omega^w_{\gamma_\alpha}] + \delta(\alpha) \\
&\le e^{\epsilon(\alpha)}\Pr[\mathcal{M}(\boldsymbol{\pi}') \in A \cap \Omega^w_{\gamma_\alpha}] + \delta(\alpha) \\
&\le e^{\epsilon(\alpha)}\Pr[\mathcal{M}(\boldsymbol{\pi}') \in A] + \delta(\alpha). \qquad (109)
\end{aligned}$$

Hence, $\mathcal{M}$ satisfies $(\epsilon(\alpha), \delta(\alpha))$-differential privacy on $\Omega^\pi_{\gamma_\alpha}$. □

Proof of Corollary 3 To prove (i), recall equation 96 and equation 97 from the proof of Theorem 2. We have

$$|\mathcal{I}_1| \le 2n\psi'(\zeta), \qquad (110)$$

where $\zeta \in [\alpha/N, \alpha/N + n]$. Using the asymptotic expansion of the trigamma function,

$$\psi'(\zeta) = \frac{1}{\zeta} + \mathcal{O}\left(\frac{1}{\zeta^2}\right), \qquad (111)$$

and combining equation 110 and equation 111, we obtain

$$|\mathcal{I}_1| \le 2n\left[\frac{1}{\zeta} + \mathcal{O}\left(\frac{1}{\zeta^2}\right)\right]. \qquad (112)$$

Since $\zeta \ge \alpha/N \to \infty$ as $\alpha \to \infty$, the right-hand side converges to zero. Therefore,

$$\mathcal{I}_1 \to 0 \qquad (113)$$

as $\alpha \to \infty$.

To prove (ii), recall the input and output sets from Theorem 2:

$$\Omega^{\pi}_{\gamma_\alpha} = \{\boldsymbol{\pi} \in \Delta_N : \pi_j \geq \gamma_\alpha,\ j \in [\![N]\!]\}\,, \tag{114}$$

$$\Omega^{w}_{\gamma_\alpha} = \{\boldsymbol{w} \in \Delta_N : w_j \geq \gamma_\alpha,\ j \in [\![N]\!]\}\,. \tag{115}$$

By the definition of $\delta(\alpha)$,

$$\delta(\alpha) = \sup_{\boldsymbol{\pi} \in \Omega^{\pi}_{\gamma_\alpha}} \Pr\big[\mathcal{M}(\boldsymbol{\pi}) \in \big(\Omega^{w}_{\gamma_\alpha}\big)^c\big]\,. \tag{116}$$

For any $\boldsymbol{\pi} \in \Omega^{\pi}_{\gamma_\alpha}$, the union bound gives

$$\begin{aligned}\Pr\big[\mathcal{M}(\boldsymbol{\pi}) \in \big(\Omega^{w}_{\gamma_\alpha}\big)^c\big] &= \Pr\big[(w_1, \ldots, w_N) \in \big(\Omega^{w}_{\gamma_\alpha}\big)^c \,\big|\, z_{1:n}(s)\big] \\ &= \Pr\left[\bigcup_{j=1}^{N} \{w_j < \gamma_\alpha\} \,\middle|\, z_{1:n}(s)\right] \\ &\leq \sum_{j=1}^{N} \Pr[w_j < \gamma_\alpha \,|\, z_{1:n}(s)]\,. \end{aligned} \tag{117}$$

By the properties of the Dirichlet distribution,

$$\mathbb{E}[w_j \mid z_{1:n}(s)] = \frac{\alpha/N + n\pi_j(s)}{\alpha + n}, \tag{118}$$

and

$$\begin{aligned}\mathrm{Var}[w_j \mid z_{1:n}(s)] &= \frac{\mathbb{E}[w_j \mid z_{1:n}(s)]\,(1 - \mathbb{E}[w_j \mid z_{1:n}(s)])}{\alpha + n + 1} \\ &\leq \frac{1}{4(\alpha + n + 1)}, \end{aligned} \tag{119}$$

where the inequality follows from $\mathbb{E}[w_j \mid z_{1:n}(s)] \in [0, 1]$ and $x(1-x) \leq 1/4$ for $x \in [0, 1]$.

We first show that $\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha$ is positive for sufficiently large α. By the definition $\gamma_\alpha = [1 - (1+\alpha)^{-q}]/N$ and equation 118,

$$\begin{aligned}\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha &= \left(\mathbb{E}[w_j \mid z_{1:n}(s)] - \frac{1}{N}\right) + \left(\frac{1}{N} - \gamma_\alpha\right) \\ &= \frac{N n_j(s) - n}{N(\alpha + n)} + \frac{(1+\alpha)^{-q}}{N}. \end{aligned} \tag{120}$$

Since $N n_j(s) - n \geq -n$, it follows that

$$\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha \geq \frac{1}{N}\left[(1+\alpha)^{-q} - \frac{n}{\alpha + n}\right]. \tag{121}$$

To show that the term in brackets is positive for sufficiently large α, consider

$$\begin{aligned}\frac{\dfrac{n}{\alpha + n}}{(1+\alpha)^{-q}} &= \frac{n(1+\alpha)^q}{\alpha + n} \\ &= n\alpha^{q-1}\frac{(1 + 1/\alpha)^q}{1 + n/\alpha}. \end{aligned} \tag{122}$$

For fixed n, as $\alpha \to \infty$,

$$\alpha^{q-1} \to 0, \qquad \frac{(1+1/\alpha)^q}{1 + n/\alpha} \to 1, \tag{123}$$

because $0 < q < 1/2 < 1$. Hence,

$$\frac{\dfrac{n}{\alpha + n}}{(1+\alpha)^{-q}} \to 0. \tag{124}$$

Therefore, there exists $\alpha_0 > 0$ such that, for every $\alpha \geq \alpha_0$,

$$\frac{n}{\alpha + n} < (1+\alpha)^{-q}. \tag{125}$$

Combining equation 121 and equation 125 gives

$$\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha > 0 \tag{126}$$

for all $\alpha \geq \alpha_0$.

We can now apply Chebyshev's inequality. For each j and $\alpha \geq \alpha_0$,

$$\Pr[w_j < \gamma_\alpha \,|\, z_{1:n}(s)] \leq \Pr[|w_j - \mathbb{E}[w_j \mid z_{1:n}(s)]| \geq \mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha \,|\, z_{1:n}(s)] \tag{127a}$$

$$\leq \frac{\mathrm{Var}[w_j \mid z_{1:n}(s)]}{\left(\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha\right)^2} \tag{127b}$$

$$\leq \frac{1}{4(\alpha+n+1)\left(\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha\right)^2}. \tag{127c}$$

Since the right-hand side of equation 121 is positive for $\alpha \geq \alpha_0$, squaring both sides gives

$$\left(\mathbb{E}[w_j \mid z_{1:n}(s)] - \gamma_\alpha\right)^2 \geq \frac{1}{N^2}\left[(1+\alpha)^{-q} - \frac{n}{\alpha+n}\right]^2. \tag{128}$$

Combining equation 127 and equation 128 therefore yields

$$\Pr[w_j < \gamma_\alpha \,|\, z_{1:n}(s)] \leq \frac{N^2}{4(\alpha+n+1)\left[(1+\alpha)^{-q} - \frac{n}{\alpha+n}\right]^2}. \tag{129}$$

Applying equation 129 in equation 117 gives

$$\Pr\left[\mathcal{M}(\boldsymbol{\pi}) \in \left(\Omega^w_{\gamma_\alpha}\right)^c\right] \leq \frac{N^3}{4(\alpha+n+1)\left[(1+\alpha)^{-q} - \frac{n}{\alpha+n}\right]^2}. \tag{130}$$

It remains to show that the right-hand side converges to zero. From equation 122,

$$\frac{\frac{n}{\alpha+n}}{(1+\alpha)^{-q}} \to 0. \tag{131}$$

Therefore,

$$\frac{(1+\alpha)^{-q} - \frac{n}{\alpha+n}}{(1+\alpha)^{-q}} = 1 - \frac{\frac{n}{\alpha+n}}{(1+\alpha)^{-q}} \to 1. \tag{132}$$

Hence,

$$\left[(1+\alpha)^{-q} - \frac{n}{\alpha+n}\right]^2 \sim (1+\alpha)^{-2q}. \tag{133}$$

It follows that

$$\begin{aligned}(\alpha+n+1)\left[(1+\alpha)^{-q} - \frac{n}{\alpha+n}\right]^2 &\sim (\alpha+n+1)(1+\alpha)^{-2q} \\ &\sim \alpha^{1-2q}.\end{aligned} \tag{134}$$

Since $0 < q < 1/2$, we have $1 - 2q > 0$, and therefore

$$\alpha^{1-2q} \to \infty. \tag{135}$$

Thus, the right-hand side of equation 130 converges to zero.

Finally, since this upper bound is independent of the input $\boldsymbol{\pi}$, taking the supremum over $\boldsymbol{\pi} \in \Omega^{\pi}_{\gamma_\alpha}$ and using equation 116 gives

$$\delta(\alpha) \to 0. \tag{136}$$

This completes the proof of (ii). □

Proof of Remark 1 Since $0 \le w_j \le 1$,

$$\mathbb{E}[w_j \mid z_{1:n}(s)] = \int_0^1 \Pr(w_j > t \mid z_{1:n}(s))\, dt.$$

By Theorem 2, for every $t \in [0,1]$,

$$\Pr(w_j > t \mid z_{1:n}(s)) \le e^{\epsilon(\alpha)} \Pr(w_j > t \mid z'_{1:n}(s)) + \delta(\alpha).$$

Integrating over $t \in [0,1]$ gives

$$\mathbb{E}[w_j \mid z_{1:n}(s)] \le e^{\epsilon(\alpha)} \mathbb{E}[w_j \mid z'_{1:n}(s)] + \delta(\alpha).$$

The reverse inequality follows by interchanging the neighboring datasets. □

A.3 Threat Model and Adversary Knowledge

We consider an adversary who has access to the released synthetic data generated from the posterior predictive distribution and seeks to reconstruct the underlying centroid trajectories. The proposed mechanism and its general hyperparameter settings, including the mean-reversion rate β_j and the functional specification of the diffusion coefficient $\sigma_j(s)$, may be known to the adversary. Thus, the proposed trajectory protection does not rely on hiding the mechanism or its general hyperparameter settings.

The original observations and data-dependent quantities used internally by the data owner are not released. In particular, when the target means m_j are constructed from empirical cluster information, as in the experiments of Section 5, their instantiated values remain internal to the data owner. This distinction is important because releasing such data-dependent target means could itself reveal information about the underlying empirical centroids.

For centroid trajectories, we consider a reconstruction adversary that may repeatedly access synthetic releases, cluster the released samples, and apply temporal filtering to estimate the underlying trajectories. Our experiments instantiate this adversary through repeated K-means clustering followed by Savitzky-Golay filtering. The resulting evaluation examines whether the underlying temporal patterns can be recovered from the released synthetic data. We do not claim that $U_s(\beta_j)$-diffusion privacy provides an adversary-independent indistinguishability guarantee; rather, it provides the trajectory protection characterized by Theorem 1 and evaluated against the specified reconstruction attack.

This threat model is distinct from that for mixture weights. For the weights, Theorem 2 provides a formal $(\epsilon(\alpha), \delta(\alpha))$-DP guarantee with respect to neighboring datasets, which does not rely on restricting the adversary's knowledge of the mechanism.

A.4 Additional Experimental Results

This section presents additional numerical results that further illustrate the behavior of the proposed procedure and provide a more comprehensive evaluation of its performance.

A.4.1 Adaptive Truncation Level N

This section provides numerical illustrations to guide the selection of the adaptive truncation level N. We first describe the adaptive truncation rule used to determine N.

Remark 3 *To determine the truncation level N in the finite DirPMM, we adopt the adaptive rule of Zarepour & Al-Labadi (2012),*

$$N = \inf\left\{ \ell : \frac{E_{\ell,\ell}}{\sum_{i=1}^{\ell} E_{i,\ell}} < \rho \right\},$$

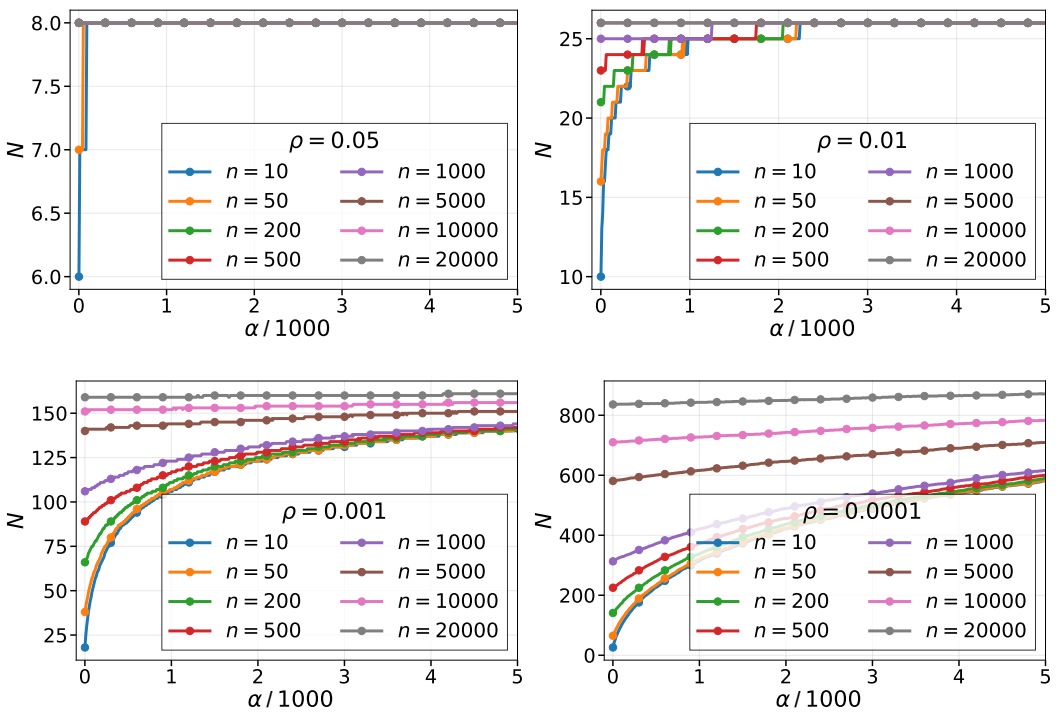

Figure 6: Adaptive truncation level N as a function of α and n for different thresholds ρ. From top left to bottom right, $\rho \in \{0.05, 0.01, 0.001, 0.0001\}$.

where $E_{i,\ell} \overset{\text{i.i.d.}}{\sim} \text{Gamma}((\alpha+n)/\ell, 1)$ *and* $\rho \in (0,1)$ *is a prescribed threshold. The truncation is selected when the next normalized weight falls below* ρ*.*

Following Remark 3, N depends on the threshold ρ, the concentration parameter α, and the sample size n. Figure 6 illustrates the resulting values of N for different choices of these parameters.

Several patterns can be observed. For a fixed threshold ρ, the adaptive truncation level N increases with α and n initially, but eventually converges to a stable value as both parameters become sufficiently large. In contrast, for fixed α and n, decreasing the threshold ρ results in a larger truncation level N, since more mixture components are retained.

In our experiments, we first specify the desired threshold ρ and concentration parameter α, and then determine the adaptive truncation level N according to the sample size of the target dataset. For example, in the toy example of Section 5 in the main paper, $n = 50$, $\alpha = 10^{-5}$, and $\rho = 0.05$ yield $N = 7$.

A.4.2 Functional Behavior of $\epsilon(\alpha)$ and $\delta(\alpha)$

For the toy example in the main paper, we further examine the $(\epsilon(\alpha), \delta(\alpha))$-DP guarantee. We consider different sample sizes $n \in \{50, 500, 1000\}$ and different choices of q. For each α, the number of latent groups N is selected using the adaptive truncation rule with $\rho = 0.05$, consistent with the experimental setup. As established in Theorem 2, $\epsilon(\alpha)$ is independent of s, whereas $\delta(\alpha)$ depends on the time-varying proportions $\pi_j(s) = n_j(s)/n$. We therefore report the time-averaged $\delta(\alpha)$, obtained by averaging its values over s. Figure 7 shows that both $\epsilon(\alpha)$ and $\delta(\alpha)$ decrease toward zero as α increases, as established in Corollary 3. The decreasing behavior of $\epsilon(\alpha)$ is similar across the considered values of q, whereas $\delta(\alpha)$ decreases substantially as q decreases. We therefore set $q = 0.1$ when reporting the $(\epsilon(\alpha), \delta(\alpha))$-DP guarantee for mixture weights in our experiments.

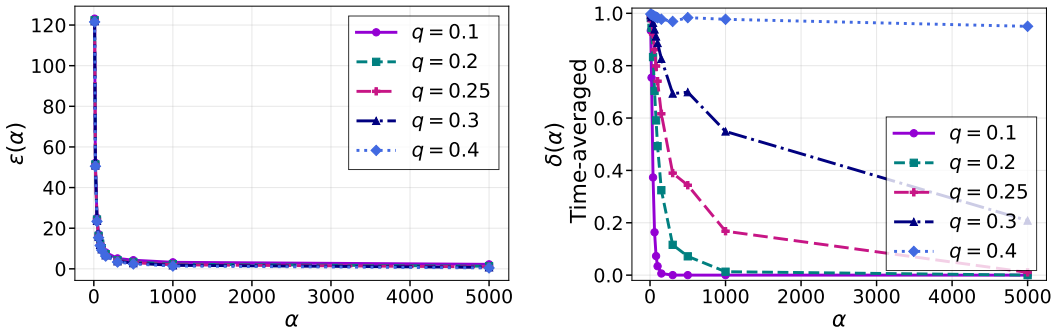

(a) $n = 50$

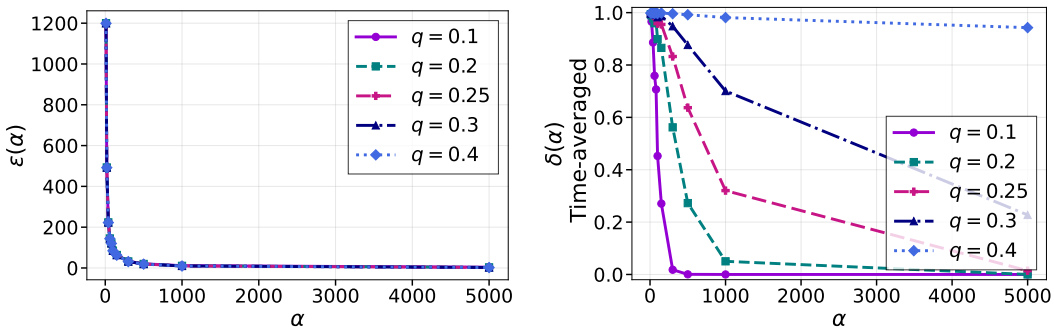

(b) $n = 500$

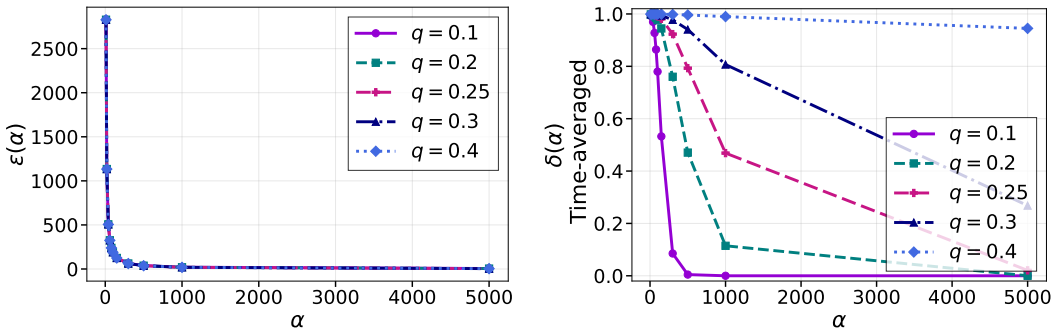

(c) $n = 1000$

Figure 7: Effect of α on $\epsilon(\alpha)$ (left) and $\delta(\alpha)$ (right) across different values of q, with $\rho = 0.05$. Each row corresponds to a different sample size n.

A.4.3 Toy Example with A More Complex Periodic Diffusion Coefficient

Recall the toy example used in the main paper, consisting of a two-modal Gaussian population with $d = 1$ over $s \in [0, 20]$:

$$X(s) \sim 0.75\,\mathrm{N}(\mu_{\mathrm{true},1}(s), 0.16^2) + 0.25\,\mathrm{N}(\mu_{\mathrm{true},2}(s), 0.16^2),$$

where $\mu_{\mathrm{true},1}(s)$ and $\mu_{\mathrm{true},2}(s)$ are the true centroids of groups 1 and 2, respectively. Let

$$h_{[a,b]}(s) = \{\tanh((s-a)/0.08) - \tanh((s-b)/0.08)\}/2,$$

which is approximately one on $[a, b]$ and zero outside. We set

$$\begin{aligned}\mu_{\mathrm{true},1}(s) &= 4 + 0.045\sin(4s) + 0.025\sin(8s + 0.4) - 1.25h_{[4,6]}(s)(1 + 0.02\sin(10s)) \\ &\quad + 0.85h_{[9,12]}(s)(1 + 0.025\sin(9s)), \\ \mu_{\mathrm{true},2}(s) &= 0.04\sin(4.3s + 0.5) + 0.02\sin(8.5s) + 1.55h_{[9,12]}(s)(1 + 0.025\sin(9.5s + 0.5)).\end{aligned}$$

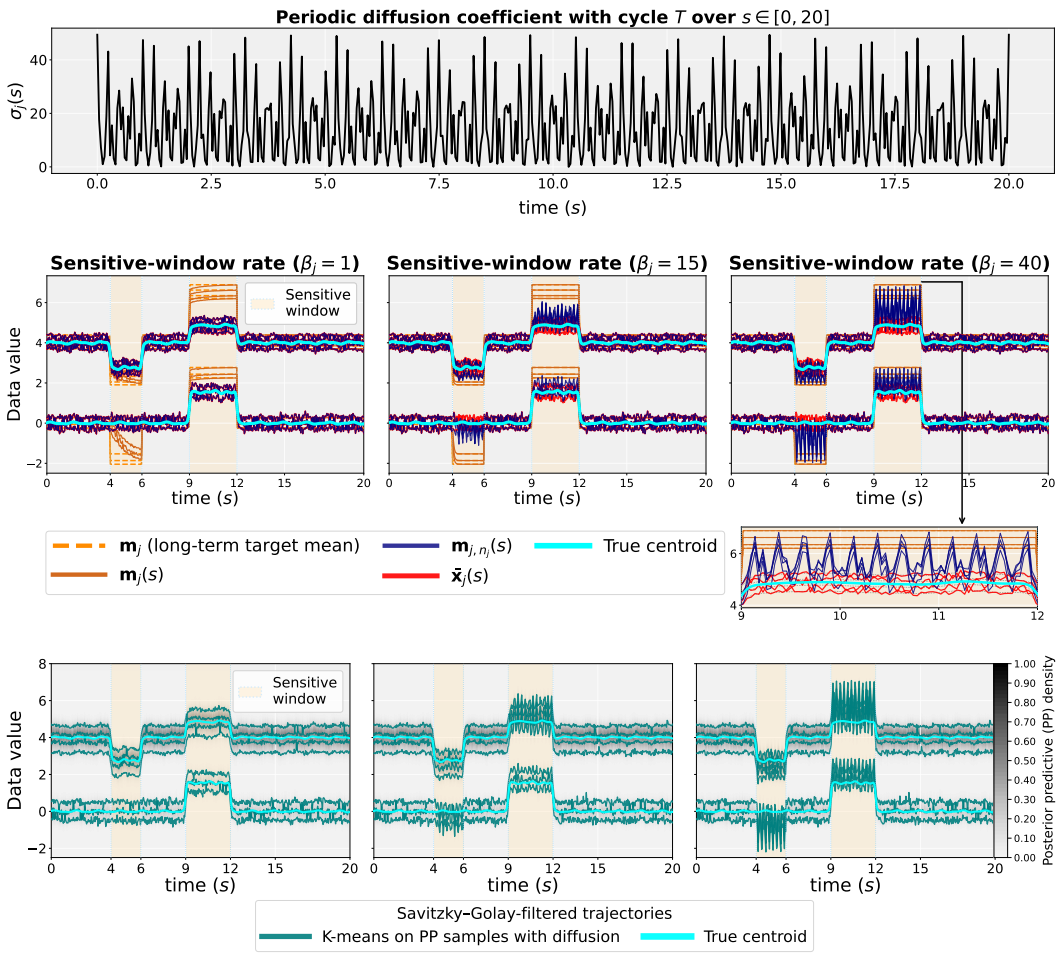

Figure 8: Effect of the complex periodic diffusion coefficient on centroid trajectory protection. **Top:** Periodic diffusion coefficient $\sigma_j(s)$. **middle:** Posterior mean trajectories under sensitive-window rates $\beta_j \in \{1, 15, 40\}$. **Bottom:** Filtered trajectories from PP samples with and without diffusion, compared with the true centroids.

To further assess the effect of the functional form of the diffusion coefficient on the centroid trajectory protection, we now consider a more complex periodic specification that combines a nonlinear envelope with an oscillatory component. Let T denote the cycle length and define the local time within the kth cycle as

$$\tau(s) = s - (s_0 + kT), \qquad s \in [s_0 + kT, s_0 + (k+1)T), \qquad k \in \mathbb{N}_0,$$

so that $\tau(s) \in [0, T)$ within each cycle. The diffusion coefficient is defined as

$$\sigma_j(s) = \begin{cases} \sigma_j \left(1 - \dfrac{2\tau(s)}{T}\right)^p \left|\sin\left(\dfrac{4\pi\tau(s)}{T} + \phi_D\right)\right|, & 0 \le \tau(s) < \dfrac{T}{2}, \\ \sigma_j \left(\dfrac{2\tau(s)}{T} - 1\right)^p \left|\sin\left(2\pi\left[\dfrac{2\tau(s)}{T} - 1\right] + \phi_I\right)\right|, & \dfrac{T}{2} \le \tau(s) < T, \end{cases}$$

where $\sigma_j > 0$ controls the overall magnitude of the diffusion for the jth cluster. In this experiment, we use the same diffusion magnitude for all clusters, i.e., $\sigma_j = \sigma = 50$ for all $j \in N$. The parameter $p > 0$ controls the shape of the envelope, and ϕ_D and ϕ_I are phase parameters for the decreasing and increasing portions of each cycle, respectively.

The parameter p determines the curvature of the envelope. When $p = 1$, the envelope changes linearly within each half-cycle. For $p > 1$, the envelope is nonlinear and more strongly curved: during the decreasing half-cycle it initially decreases more slowly and then drops more rapidly as $\tau(s)$ approaches $T/2$, while during the increasing half-cycle it initially increases slowly and becomes steeper toward the end of the cycle. In contrast, for $0 < p < 1$, the decreasing envelope drops more

rapidly near the beginning of the half-cycle and then flattens, while the increasing envelope rises rapidly near $T/2$ and subsequently flattens.

In our experiment, we set

$$T = 0.25, \qquad p = 1.2, \qquad \phi_D = \phi_I = 5000.$$

Thus, $p = 1.2$ introduces a moderate nonlinear curvature relative to the linear case $p = 1$, while the absolute sinusoidal factor introduces within-cycle oscillations. This construction provides a more complex time-varying diffusion pattern for examining the sensitivity of the resulting posterior mean to the specification of $\sigma_j(s)$. The form of $\sigma_j(s)$ is illustrated in the top panel of Figure 8.

The middle panels of Figure 8 show the posterior mean trajectories $m_{j,n_j}(s)$ obtained using this more complex diffusion coefficient. We consider three sensitive-window rates, with the same rate used for all $j \in N$: $\beta_j \in \{1, 15, 40\}$. As β_j increases, the effect of the diffusion becomes more visible within the sensitive windows, and the posterior means show stronger and more frequent fluctuations between $\bar{x}_j$ and $m_j(s)$. In particular, the zoomed-in region in Figure 8 clearly shows the repeated local fluctuations induced by the oscillatory form of $\sigma_j(s)$. Compared with the simpler periodic $\sigma_j(s)$ used in the main experimental results, this more complex specification produces a more irregular perturbation pattern in the posterior mean while retaining its time-varying behavior. This example therefore highlights the flexibility of the proposed diffusion construction in producing different perturbation patterns for centroid trajectories, beyond the simple form of $\sigma_j(s)$ considered in the main experiments.

We next evaluate whether a third party, having access only to the released synthetic data, can reconstruct the true centroid trajectories through repeated clustering and filtering. The bottom panels of Figure 8 examine this filtering-based reconstruction. At every time point s, we repeat the PP sampling and clustering 100 times. In each repetition, we generate 50 synthetic samples from the PP distribution and apply K-means clustering to estimate the centroids. Consistent with our DirPMM setting, we set $K = 7$. Thus, for each cluster at each time point s, we obtain 100 estimated centroid values. These estimates are first averaged at each time point to form a single averaged centroid trajectory. The Savitzky-Golay filter (Savitzky & Golay, 1964) is then applied to the averaged trajectory. The filter fits a polynomial by least squares to neighboring points along the averaged trajectory and uses the fitted values to obtain a smoothed centroid trajectory. The resulting filtered trajectory, shown in the bottom panels of Figure 8, is used to assess whether this filtering can remove the diffusion-induced perturbations and reveal the true centroid trajectory. The results show a clear dependence on β_j. For smaller values of β_j, the reconstructed paths remain relatively close to the true centroids. As β_j increases, the reconstructed paths deviate more strongly within the sensitive windows and show jumps of nearly similar size across clusters. This makes it difficult for a third party to identify which cluster undergoes the actual jump, which may itself be sensitive information. Hence, increasing β_j makes it increasingly difficult to infer the original trajectory structure from the released PP samples.

Similar to the results in the main paper, Figure 9 shows a clear privacy-utility trade-off induced by the diffusion trajectories. The top panels show $\mathcal{U}_s(\beta_j)$ across time. For $\beta_j = 1$, the values remain close to one, indicating little protection within the sensitive windows. As β_j increases to 15 and 40, $\mathcal{U}_s(\beta_j)$ decreases substantially within these windows, indicating stronger protection of the centroid trajectories. This improved protection comes with some loss in utility. As shown in the bottom panels, the MMD between the PP and original samples increases mainly within the sensitive windows as β_j becomes larger, while remaining relatively small outside these windows. As expected, these results also show the effect of the more complex perturbation pattern on utility. Overall, stronger diffusion provides greater protection within the sensitive windows at the cost of some utility loss.

A.4.4 Membership Inference Attack Protocol

To empirically evaluate whether the differentially private mixture weights leak membership information, we consider the following membership inference attack protocol.

1. Construct two datasets,

$$\mathbf{X}_{1:n}(s), \qquad \mathbf{X}'_{1:n}(s),$$

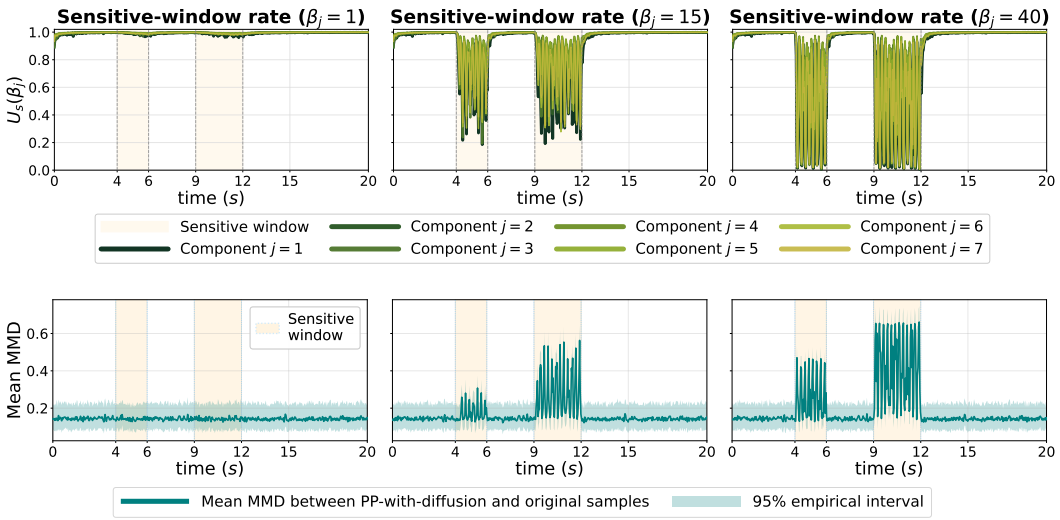

Figure 9: Privacy-utility trade-off induced by diffusion trajectories in the toy example. **Top:** $\mathcal{U}_s(\beta_j)$ across s for different β_j. **Bottom:** MMD between PP and original samples.

that differ in the membership of nm observations. Specifically, $\mathbf{X}'_{1:n}(s)$ is obtained from $\mathbf{X}_{1:n}(s)$ by replacing nm observations with observations drawn from group 2 of the toy population. The case $nm = 1$ corresponds to neighboring datasets.

2. Fit the proposed DDirPMM independently on $\mathbf{X}_{1:n}(s)$ and $\mathbf{X}'_{1:n}(s)$ over repeated runs.
3. For each run, compute the mixture-weight vectors used in the PP distributions,

$$\widehat{\mathbf{w}}(s) = \big(\mathbb{E}(\omega_1 \mid z_{1:n}(s)), \ldots, \mathbb{E}(\omega_N \mid z_{1:n}(s))\big),$$

and

$$\widehat{\mathbf{w}}'(s) = \big(\mathbb{E}(\omega_1 \mid z'_{1:n}(s)), \ldots, \mathbb{E}(\omega_N \mid z'_{1:n}(s))\big).$$

4. Assign label 0 to weight vectors obtained from $\mathbf{X}_{1:n}(s)$ and label 1 to those obtained from $\mathbf{X}'_{1:n}(s)$.
5. Train a membership attacker (e.g., Logistic Regression (Cox, 1959), Random Forest (Breiman, 2001), Support Vector Machine (Cortes & Vapnik, 1995), Multi-Layer Perceptron (Rumelhart et al., 1986), Gradient Boosting (Friedman, 2001), or the likelihood-ratio test Neyman & Pearson (1992)) on the weight vectors and their labels to predict which dataset was used to fit the model.
6. Evaluate the attacker on an independent set of weight vectors using the receiver operating characteristic (ROC) curve and its area under the curve (AUC); see Fawcett (2006).

The attacker's goal is to determine whether a weight vector comes from the model fitted on $\mathbf{X}_{1:n}(s)$ (label 0) or $\mathbf{X}'_{1:n}(s)$ (label 1). An AUC close to 0.5 (random guessing) means that the differing membership is difficult to detect from the mixture weights, indicating that its effect on group proportions is protected. A larger AUC indicates greater membership leakage through the mixture weights.

A.5 Variational Inference for the Non-Conjugate Posterior in Equation 6

Variational inference approximates the posterior distribution $\mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))$ by a parametric family $\{q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi}) : \boldsymbol{\psi} \in \Psi\}$, where the variational parameters $\boldsymbol{\psi}$ are optimized by minimizing the Kullback-Leibler (KL) divergence to the true posterior. Specifically,

$$\begin{aligned}
&\mathrm{D}_{\mathrm{KL}}\left(q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi}) \| \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))\right) \\
&= \mathbb{E}_q\left[\log q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi})\right] - \mathbb{E}_q\left[\log \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))\right] \\
&= \mathbb{E}_q\left[\log q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi})\right] - \mathbb{E}_q\left[\log \mathring{g}(\mathbf{x}_{I_j}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j)\right] \\
&\quad - \mathbb{E}_q\left[\log \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j)\right] + \log \mathring{g}(\mathbf{x}_{I_j}(s)),
\end{aligned} \tag{137}$$

where the last term follows from

$$\mathbb{E}_q\left[\log \mathring{g}(\mathbf{x}_{I_j}(s))\right] = \log \mathring{g}(\mathbf{x}_{I_j}(s)).$$

We consider the mean-field variational approximation (Jordan et al., 1999), where

$$q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi}) = q(\boldsymbol{\mu}_j(s); \boldsymbol{\psi}_{\boldsymbol{\mu}_j}) \times q(\Sigma_j; \boldsymbol{\psi}_{\Sigma_j}), \tag{138}$$

with $\boldsymbol{\psi} = (\boldsymbol{\psi}_{\boldsymbol{\mu}_j}, \boldsymbol{\psi}_{\Sigma_j}) \in \Psi$.

Substituting equation 138 into equation 137 gives

$$\begin{aligned}
&\mathrm{D}_{\mathrm{KL}}\left[q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi}) \| \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))\right] \\
&= \mathbb{E}_q\left[\log q(\boldsymbol{\mu}_j(s); \boldsymbol{\psi}_{\boldsymbol{\mu}_j})\right] + \mathbb{E}_q\left[\log q(\Sigma_j; \boldsymbol{\psi}_{\Sigma_j})\right] \\
&\quad - \mathbb{E}_q\left[\log \mathring{g}(\mathbf{x}_{I_j}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j)\right] \\
&\quad - \mathbb{E}_q\left[\log \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j)\right] + \log \mathring{g}(\mathbf{x}_{I_j}(s)).
\end{aligned} \tag{139}$$

Since the final term is independent of q, minimizing the KL divergence is equivalent to maximizing the evidence lower bound (ELBO):

$$\begin{aligned}
&\arg\min_{\Psi} \mathrm{D}_{\mathrm{KL}}\left[q(\boldsymbol{\mu}_j(s), \Sigma_j; \boldsymbol{\psi}) \| \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j \mid \mathbf{x}_{I_j}(s))\right] \\
&= \arg\max_{\Psi} \Big\{ \mathbb{E}_q\left[\log \mathring{g}(\mathbf{x}_{I_j}(s) \mid \boldsymbol{\mu}_j(s), \Sigma_j)\right] + \mathbb{E}_q\left[\log \mathring{g}(\boldsymbol{\mu}_j(s), \Sigma_j)\right] \\
&\qquad - \mathbb{E}_q\left[\log q(\boldsymbol{\mu}_j(s); \boldsymbol{\psi}_{\boldsymbol{\mu}_j})\right] - \mathbb{E}_q\left[\log q(\Sigma_j; \boldsymbol{\psi}_{\Sigma_j})\right] \Big\}.
\end{aligned} \tag{140}$$

Traditional variational inference typically assumes predefined families of distributions for the marginals in equation 138. Such restrictive choices may limit the ability to accurately approximate the true posterior distribution. To overcome this limitation, two parametric functions represented by neural networks are considered for generating $\boldsymbol{\mu}_j(s)$ and Σ_j, and the corresponding marginal distributions in equation 138 are constructed using normalizing flows (Rezende & Mohamed, 2015). The complete procedure is summarized in Algorithm 1.

A.6 Algorithms

A.7 Limitations and Future Work

Two natural extensions of our framework remain open, each addressing a distinct gap in the privacy guarantees we currently provide.

A.7.1 Diffusion Privacy for Mixture Weights

Our framework focuses on protecting centroid trajectories, but in some dynamic settings the mixture weights themselves carry sensitive information. A concrete example is the use of interacting GMMs as a proxy for interacting LLM agents (Wang et al., 2025), where weights encode evolving agent beliefs and their trajectories can reveal behavioral patterns over time.

Extending diffusion-based protection to weights is non-trivial. Centroids live in Euclidean space and admit the MR-OUP as a natural prior, but weights lie on the probability simplex and arise through Dirichlet distributions, a fundamentally different geometry. Directly applying the MR-OUP is therefore not appropriate. What is needed is a diffusion process that evolves on the simplex while preserving Dirichlet structure, so that the computational and inferential advantages of our current framework carry over. Existing results on Dirichlet diffusion processes (Jenkins & Spanò, 2017; Bakosi & Ristorcelli, 2013) suggest this is a tractable direction, and we plan to pursue it as part of a unified treatment of trajectory privacy for all model parameters.

A.7.2 Differential Privacy for Centroid Trajectories

Our diffusion mechanism protects centroid trajectories from reconstruction but does not provide differential privacy in the classical sense. For applications requiring both forms of protection, two directions seem promising.

The simpler one is to privatize the long-term target means $\mathbf{m}_j$ of the MR-OUPs directly, by adding calibrated DP noise before inference. Since $\mathbf{m}_j$ governs where each centroid trajectory is drawn toward over time, perturbing it at this level could naturally extend DP guarantees to the trajectory without requiring changes to the rest of the model.

A more ambitious direction follows from a preliminary observation in our analysis: under suitable regularity conditions, the posterior centroid trajectories appear to admit a Gaussian process representation. If this can be made rigorous, it would open a connection to the growing literature on differentially private GPs (Honkela & Melkas, 2021), potentially yielding trajectory-level DP guarantees with well-understood privacy-utility trade-offs. This direction requires substantially more theoretical work, but the payoff, a principled and compositional privacy framework for the full model, would be significant.

We intend to pursue both directions and evaluate them under a shared DP budget, with the goal of identifying which offers better utility for a given privacy level.

Algorithm 1: Neural Flow Approximation of $q_j(\boldsymbol{\mu}_j, \Sigma_j)$

Input: Component data $\mathbf{x}_{I_j} = \{\mathbf{x}_i : z_i = j\}$; prior density $\mathring{g}(\cdot, \cdot)$; number of Monte Carlo draws n_{MC}; numerical constant $\epsilon_0 = 10^{-6}$; Adam learning rate $\rho = 10^{-3}$; Adam parameters $(\varepsilon = 10^{-8}, \beta_1 = 0.9, \beta_2 = 0.999)$.

Output: Optimized neural variational parameters $\boldsymbol{\vartheta}_j^* = (\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}^*, \boldsymbol{\vartheta}_{\Sigma_j}^*)$.

Initialize neural flow parameters $\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}$ and $\boldsymbol{\vartheta}_{\Sigma_j}$;

repeat

  Set $\widehat{\mathcal{L}}_j \leftarrow 0$;

  for $\ell = 1$ *to* n_{MC} **do**

    Sample noise variables

$$\boldsymbol{\epsilon}_{\boldsymbol{\mu}_j}^{(\ell)} \sim \mathrm{N}(\mathbf{0}, I), \qquad \boldsymbol{\epsilon}_{\Sigma_j}^{(\ell)} \sim \mathrm{N}(\mathbf{0}, I).$$

    Generate neural-flow outputs

$$\boldsymbol{\mu}_j^{(\ell)} \leftarrow f_{\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}}(\boldsymbol{\epsilon}_{\boldsymbol{\mu}_j}^{(\ell)}), \qquad \mathbf{r}_j^{(\ell)} \leftarrow h_{\boldsymbol{\vartheta}_{\Sigma_j}}(\boldsymbol{\epsilon}_{\Sigma_j}^{(\ell)}), \qquad \mathbf{r}_j^{(\ell)} \in \mathbb{R}^{d(d+1)/2}.$$

    ; `// $\mathbf{r}_j^{(\ell)}$ stores the free entries of a Cholesky factor`

    Let $\mathcal{I}(k, \ell)$ denote the fixed index of entry (k, ℓ), $1 \le \ell \le k \le d$, in the vector $\mathbf{r}_j^{(\ell)}$;

    Define

$$L_j^{(\ell)} = \left[L_{j,km}^{(\ell)}\right]_{k,m=1}^{d} \in \mathbb{R}^{d \times d},$$

    where

$$L_{j,km}^{(\ell)} = \begin{cases} \mathrm{softplus}\left(r_{j,\mathcal{I}(k,k)}^{(\ell)}\right) + \epsilon_0, & k = m, \\ r_{j,\mathcal{I}(k,m)}^{(\ell)}, & k > m, \\ 0, & k < m. \end{cases}$$

    ; `// softplus$(x) = \log(1 + e^x)$ ensures positive diagonal entries`

    Set

$$\Sigma_j^{(\ell)} \leftarrow L_j^{(\ell)} L_j^{(\ell)\top}.$$

    ; `// $\epsilon_0 > 0$ guarantees $\Sigma_j^{(\ell)} \succ 0$`

    Compute flow log-densities

$$\log q(\boldsymbol{\mu}_j^{(\ell)}; \boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}) = \log p(\boldsymbol{\epsilon}_{\boldsymbol{\mu}_j}^{(\ell)}) - \log \left| \det J_{f_{\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}}}(\boldsymbol{\epsilon}_{\boldsymbol{\mu}_j}^{(\ell)}) \right|,$$

$$\log q(\Sigma_j^{(\ell)}; \boldsymbol{\vartheta}_{\Sigma_j}) = \log p(\boldsymbol{\epsilon}_{\Sigma_j}^{(\ell)}) - \log \left| \det J_{h_{\boldsymbol{\vartheta}_{\Sigma_j}}}(\boldsymbol{\epsilon}_{\Sigma_j}^{(\ell)}) \right|.$$

    ; `// Change-of-variables formula for normalizing flows (Rezende & Mohamed, 2015)`

    Update the Monte Carlo estimate of the ELBO:

$$\widehat{\mathcal{L}}_j \leftarrow \widehat{\mathcal{L}}_j + \frac{1}{n_{\mathrm{MC}}} \Big[\log \mathring{g}(\mathbf{x}_{I_j} \mid \boldsymbol{\mu}_j^{(\ell)}, \Sigma_j^{(\ell)}) + \log \mathring{g}(\boldsymbol{\mu}_j^{(\ell)}, \Sigma_j^{(\ell)}) - \log q(\boldsymbol{\mu}_j^{(\ell)}; \boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}) - \log q(\Sigma_j^{(\ell)}; \boldsymbol{\vartheta}_{\Sigma_j}) \Big].$$

  Update neural variational parameters:

$$\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j} \leftarrow \mathrm{Adam}\left(\nabla_{\boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}} \widehat{\mathcal{L}}_j, \boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}, \rho, \varepsilon, \beta_1, \beta_2\right),$$

$$\boldsymbol{\vartheta}_{\Sigma_j} \leftarrow \mathrm{Adam}\left(\nabla_{\boldsymbol{\vartheta}_{\Sigma_j}} \widehat{\mathcal{L}}_j, \boldsymbol{\vartheta}_{\Sigma_j}, \rho, \varepsilon, \beta_1, \beta_2\right).$$

until *convergence*;

return $q_j(\cdot, \cdot; \boldsymbol{\vartheta}_j^*) = q(\cdot; \boldsymbol{\vartheta}_{\boldsymbol{\mu}_j}^*) q(\cdot; \boldsymbol{\vartheta}_{\Sigma_j}^*)$;