Sequential Membership Inference Attacks

Thomas Michel
Univ. Lille, Inria, CNRS, Centrale Lille
UMR 9189-CRIStAL, France

Debabrota Basu
Univ. Lille, Inria, CNRS, Centrale Lille
UMR 9189-CRIStAL, France

Emilie Kaufmann
Univ. Lille, CNRS, Inria, Centrale Lille
UMR 9189-CRIStAL, France

Abstract

Modern AI models are not static. They go through multiple updates in their lifecycles. We propose to design Sequential Membership Inference (SeMI) attacks leading to tighter privacy audits by exploiting the sequence of models and injecting a target canary at a controlled insertion time. First, for empirical mean computation, we develop SeMI^*, an optimal SeMI attack to identify the presence of a target inserted at a specific insertion step. We derive the power of SeMI^* to show that accessing the model sequence yields more powerful MI attacks than scrutinising only the final model. SeMI^* exhibits an isolation property. Its power depends on the statistics obtained right before and after insertion of the target. Leveraging this insight, we develop practical white-box (accessing model gradients) and black-box (accessing loss) SeMI attacks against models trained with (DP-)SGD. Across datasets and models trained with (DP-)SGD, our experiments show that SeMI attacks achieve higher power than snapshot-independent baselines, and yield tighter privacy audits thanks to (a) control over the insertion time and (b) observations across the model sequence. Code is available at `https://github.com/Thomick/pysemi`.

1 Introduction

Machine Learning (ML) models memorize training data, creating privacy risks for individuals whose data were used to develop them (Shokri et al., 2017; Carlini et al., 2022a). Research further establishes that even simple statistical models, like empirical mean estimation, leak membership information (Homer et al., 2008; Azize and Basu, 2025). While differential privacy (Dwork et al., 2006; Dwork and Roth, 2014) aims to bound privacy leakage of algorithms, privacy auditing quantifies these risks by testing whether specific records, referred to as target or *canary*, can be detected in the training set. *Membership Inference (MI) attacks* serve as the primary tool for auditing: an attacker, who distinguishes members from non-members with high confidence, establishes an empirical lower bound on the model's privacy leakage (Jayaraman and Evans, 2019; Jagielski et al., 2020; Nasr et al., 2023). MI attacks with higher distinguishing power yield tighter audits. In addition, the target selection matters as points 'far' from the data distribution leak more information than typical points (Azize and Basu, 2025; Carlini et al., 2022b).

Jayaraman and Evans (2019) report 5–10$\times$ gaps between theoretical DP guarantees and empirical privacy estimates. Techniques for tight auditing (Nasr et al., 2023; Steinke et al., 2023) and debugging (Tramer et al., 2022) reduce this gap by carefully crafting adversarial targets and observing per-step gradients during training. While training access is confined to the developer, downstream auditors only see the released models, whether through full weight access or query-only APIs. In addition, *modern ML models rarely exist as isolated snapshots.* They are trained or fine-tuned in

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

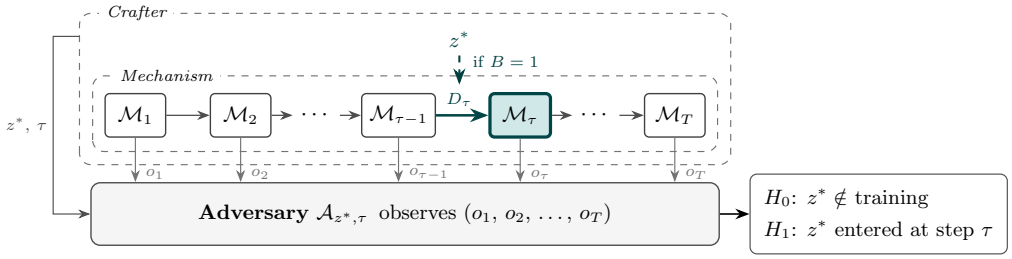

Figure 1: **Sequential MI game:** Crafter controls a mechanism that releases T outputs $(o_1, \ldots, o_T)$. With probability $1/2$, the target z^* is inserted into batch D_τ at step τ. Otherwise, the mechanism runs unmodified. The adversary observes all outputs and decides whether z^* is inserted (H_1) or not (H_0).

successive phases, with a new snapshot released after each: versioned foundation models, task-specific fine-tunes distributed on top of public backbones, continually-updated production systems, and federated models aggregated across rounds. *We show that scrutinizing the sequence of released models yields tighter audits than a single final snapshot*, while most of the auditors target a specific model snapshot.

Jagielski et al. (2023) demonstrate that combining the MI scores across model updates improves attack accuracy. They apply standard MI attacks separately to each model snapshot, and then, aggregate the scores via difference or ratio operations. Similarly, the benefits of leveraging model snapshots are empirically shown for reconstruction attacks to recover data from online learning updates (Salem et al., 2020), machine unlearning attacks (Chen et al., 2020), and leakage analysis of snapshots in NLP (Zanella-Béguelin et al., 2020). SeqMIA (Li et al., 2024), which extends TrajectoryMIA (Liu et al., 2022) from a set of losses to ordered sequences of metrics, classifies sequences computed on distilled snapshots of a single released model (Appendix A). *Our approach differs by deriving the optimal MI attacks via likelihood ratio tests and identifying the insertion time of a canary as a new actionable lever for tighter audits.*

Specifically, *we study Membership Inference (MI) when the auditor observes a sequence of T model snapshots released across training phases rather than a single static one.* We formalize this *Sequential MI (SeMI)* setting as a hypothesis testing game between a crafter and an adversary over T outputs with a choice of the insertion time τ. The auditor, a model owner or a party allowed to insert data, runs both. As crafter it chooses the canary and its insertion time, and its adversary tests the released outputs. We show how Type-I/II errors of SeMI yield a privacy lower bound for auditing. The sequential setting offers two levers for designing tight audits: *selecting targets that leak information, and choosing when to insert them*. Specifically, we ask:

1. *Does SeMI with insertion-time control improve over a static attack on the final model?*
2. *Does theory-driven design of SeMI attacks yield stronger MI attacks against (DP-)SGD?*
3. *Do SeMI attacks with canary and insertion-time control lead to tighter privacy audits?*

Our contributions address these questions affirmatively.

(1) Optimal SeMI attack, isolation property, and knowledge of τ. Motivated by the optimal static MI attack (Azize and Basu, 2025) derived from the Neyman–Pearson lemma (Neyman and Pearson, 1933), we derive the optimal SeMI attack SeMI^* as a likelihood-ratio (LR) test parameterized by the insertion time τ. For empirical mean computation on Gaussian data, we obtain the closed-form Type-I and Type-II errors. The analysis exposes an *isolation property*: consecutive outputs reveal the batch mean at τ, so the LR depends on the data only through this batch statistic. The power of SeMI^* depends on the batch size n and the target's distance from the mean, not on τ or T. A single-observation attack on the final snapshot sees the target pooled among nT samples, and its signal dilutes as training data accumulates. We further study $\mathrm{SeMI}^{\max}$ and $\mathrm{SeMI}^{\mathrm{Unif}}$, which assume τ is unknown or uniform, and find that SeMI^* dominates both.

(2) White- and black-box SeMI attacks on (DP-)SGD. Since (DP-)SGD (Kiefer and Wolfowitz, 1952; Abadi et al., 2016) is the standard for training private ML models, we extend the SeMI attack to (DP-)SGD via a batch-gradient Gaussianity assumption. This yields $\mathrm{SeMI}^{\mathrm{SGD}}$ and shows that the isolation property holds for its per-transition log-LR. In the white-box setting, $\mathrm{SeMI}^{\mathrm{SGD}}$ takes a Mahalanobis-weighted score between the parameter update and the target gradient. In the black-box

Algorithm 1 SeMI Game

Input: Mechanisms $(\mathcal{M}_1, \dots, \mathcal{M}_T)$, data distributions $(\mathcal{D}_1, \dots, \mathcal{D}_T)$, batch size n, target z^*, insertion time τ, adversary $\mathcal{A}$
Sample $B \sim \mathcal{B}(1/2)$
for $t = 1, \dots, T$ **do**
  Sample batch $D_t \sim \mathcal{D}_t^{\otimes n}$
  if $B = 1$ and $t = \tau$ **then**
    Sample $j \sim \mathrm{Uniform}(\{1, \dots, n\})$
    Replace $D_t[j] \leftarrow z^*$
  end if
  Compute $o_t = \mathcal{M}_t\left(\bigcup_{i=1}^{t} D_i\right)$
end for
Adversary guesses $\widehat{B} = \mathcal{A}(z^*, \tau, \{o_t\})$
Output: $B, \widehat{B}$

Algorithm 2 SeMI Audit

Input: Mechanisms $(\mathcal{M}_1, \dots, \mathcal{M}_T)$, data distributions $(\mathcal{D}_1, \dots, \mathcal{D}_T)$, batch size n, target z^*, insertion time τ, adversary $\mathcal{A}$, rounds R, confidence ξ, privacy δ
for $r = 1, \dots, R$ **do**
  $B_r, \widehat{B}_r \leftarrow \mathrm{SeMIGame}(\tau, \mathcal{A})$
end for
Compute Clopper–Pearson upper bounds $\overline{\alpha}_R^{\xi}, \overline{\beta}_R^{\xi}$ at confidence level $\xi/2$ (Appendix C.1)
Output: lower bound on ε from (3) with $(\overline{\alpha}_R^{\xi}, \overline{\beta}_R^{\xi})$

setting, where the auditor observes only target losses, we propose two per-transition surrogates, Loss Diff and Loss Ratio, recovering the Back-Front and Delta heuristics of Jagielski et al. (2023) as restrictions to the inserted phase.

(3) From tighter SeMI attacks to sequential privacy audits. We formalize privacy auditing of sequential mechanisms as a testing game that encodes the auditor's control over target selection and insertion time, and turn confidence intervals on the Type-I and Type-II errors of SeMI attacks on (DP-)SGD into ε lower bounds. Auditing a softmax classifier trained with DP-SGD on Fashion-MNIST, we find that (i) $\mathrm{SeMI}^{\mathrm{SGD}}$ is the only attack whose bound stays positive as the number of released snapshots T increases, (ii) on the same pair of snapshots, the white-box $\mathrm{SeMI}^{\mathrm{SGD}}$ gives tighter bounds than the black-box Loss Diff and Loss Ratio, (iii) knowing the insertion time tightens the audit relative to its worst-case counterpart, and (iv) audit tightness depends strongly on τ, giving the auditor a lever for tighter bounds beyond target selection.

2 Auditing and Membership Inference for Sequential Mechanisms

We formalise MI attacks against sequential mechanisms and their privacy audits. Data is collected over T steps: at step $t \in [T]$, a batch D_t of n i.i.d. samples from $\mathcal{D}_t$ arrives, and a sequential mechanism $\mathcal{M}_t$ is applied to the accumulated dataset, releasing $o_t \triangleq \mathcal{M}_t(\bigcup_{i=1}^{t} D_i)$. We write $\mathcal{M}$ for the global mechanism with input $D \triangleq \bigcup_{i=1}^{T} D_i$ and output $o \triangleq (o_1, \dots, o_T)$. The mechanisms and distributions may vary over time, modelling an ML system trained or fine-tuned in successive phases with each version released to users. *We aim to assess the privacy of the global mechanism $\mathcal{M}$ while observing the released outputs of the intermediate mechanisms.*

To quantify the privacy level of a mechanism, we adopt the classical (ε, δ)-Differential Privacy (DP) definition (Dwork and Roth, 2014). A mechanism $\mathcal{M}$ satisfies (ε, δ)-DP if for all measurable sets of output S and datasets D, D' differing in one entry, and some $\varepsilon > 0$ and $\delta \in [0, 1)$,

$$\Pr[\mathcal{M}(D) \in S] \leqslant e^{\varepsilon} \Pr[\mathcal{M}(D') \in S] + \delta\,. \tag{1}$$

Given some δ, *privacy auditing amounts to finding a lower bound on ε for which (1) is satisfied.*

Two dominant approaches to auditing exist. The first estimates the smallest ε for which the hockey-stick divergence $\sup_S \max\{\Pr[\mathcal{M}(D) \in S] - e^{\varepsilon} \Pr[\mathcal{M}(D') \in S], 0\}$ is below δ (Koskela and Mohammadi, 2024; Basu and Chanda, 2026). The second uses the hypothesis-testing interpretation of DP (Wasserman and Zhou, 2010; Kairouz et al., 2015), running MI tests (that are called *attacks*) and turning their error probabilities into a lower bound on ε. The second approach also guides the design of optimal MI attacks and privacy-leaking canaries (Nasr et al., 2023; Azize and Basu, 2025), and we adopt it for sequential mechanisms.

MI attacks identify whether a target z^* is present in the input dataset D from the output of a mechanism (Carlini et al., 2022a; Ye et al., 2022; Leemann et al., 2023). A crafter designs the canary z^* and decides whether to insert it. With sequential access, we further ask whether z^* belongs to a *specific sub-dataset D_τ, where $\tau \in \mathbb{N}$ is the insertion time*. A sequential MI attack tests two hypotheses parameterized by z^*:

$\mathbf{H_0}$ (OUT) : $D \sim \mathcal{P}_0$, if for all $t \in [T]$, $D_t \sim \mathcal{D}_t^{\otimes n}$.

$\mathbf{H_1^\tau}$ (IN$_\tau$) : $D \sim \mathcal{P}_1^\tau$, if for all $t \in [T]$, $D_t \sim \mathcal{D}_t^{\otimes n}$, and given $J \sim \mathcal{U}([n])$, the J-th entry of D_τ is replaced by z^*.

A **Sequential MI (SeMI) game**, stated in Algorithm 1 and illustrated in Figure 1, takes an insertion time τ[1]. The crafter draws $B \sim \mathcal{B}(1/2)$ and samples D from $\mathcal{P}_0$ if $B = 0$ or from $\mathcal{P}_1^\tau$ otherwise. The mechanism $\mathcal{M}$ produces $o_{1:T} := (o_1, \dots, o_T)$. A **SeMI attack** is an adversary $\mathcal{A}$ mapping $(\tau, z^*, o_{1:T})$ to a guess $\widehat{B} = \mathcal{A}(\tau, z^*, o_{1:T})$ for B. The associated errors are

$$\alpha(\tau, \mathcal{A}) \triangleq \mathbb{P}_{D\sim\mathcal{P}_0}[\mathcal{A}(\tau, z^*, \mathcal{M}(D)) = 1]\,, \quad \beta(\tau, \mathcal{A}) \triangleq \mathbb{P}_{D\sim\mathcal{P}_1^\tau}[\mathcal{A}(\tau, z^*, \mathcal{M}(D)) = 0]\ .$$

A general adversary that can exploit the knowledge of τ coincides with a statistical test of

$$\mathbf{H_0}\ (\text{OUT}) : (D \sim \mathcal{P}_0) \quad \text{against} \quad \mathbf{H_1^\tau}\ (\text{IN}_\tau) : (D \sim \mathcal{P}_1^\tau) \tag{2}$$

and $\alpha(\tau, \mathcal{A})$ and $\beta(\tau, \mathcal{A})$ are respectively the type I and type II errors of this test. The following extension of (Kairouz et al., 2015) relates these two testing errors to the privacy parameters (see Appendix B).

Lemma 2.1 (Connecting Sequential MI Attack Error and DP Audits)**.** *Fix a target $z^* \in \mathcal{Z}$ and an insertion time τ. For any sequential MI attack $\mathcal{A}$, if the mechanism $\mathcal{M}$ is (ε, δ)-DP, then*

$$\alpha(\tau, \mathcal{A}) + e^\varepsilon \beta(\tau, \mathcal{A}) \geqslant 1 - \delta, \qquad \beta(\tau, \mathcal{A}) + e^\varepsilon \alpha(\tau, \mathcal{A}) \geqslant 1 - \delta.$$

2.1 From SeMI Attacks to Privacy Audits

Lemma 2.1 yields a lower bound of privacy level for any insertion time and any adversary

$$\varepsilon \geqslant \log\left(\max\left[\frac{1-\delta-\alpha(\tau,\mathcal{A})}{\beta(\tau,\mathcal{A})}, \frac{1-\delta-\beta(\tau,\mathcal{A})}{\alpha(\tau,\mathcal{A})}\right]\right). \tag{3}$$

A privacy audit (Algorithm 2) runs R independent rounds of the MI game and plugs high-confidence upper bounds on $\alpha(\tau, \mathcal{A})$ and $\beta(\tau, \mathcal{A})$ into (3). In our experiments, we use a refinement of this scheme that can be applied to any attack comparing a statistic to a threshold. Instead of fixing $\mathcal{A}$ (a threshold γ on the attack statistic) in advance, we calibrate a grid of thresholds $\gamma_1, \dots, \gamma_K$ on separate runs, compute Clopper–Pearson confidence intervals (Clopper and Pearson, 1934) at every grid point, and report the maximum lower bound across the grid with a union bound. Appendix C details this procedure and a calibration-free alternative based on the DKW inequality. A tight lower bound now requires designing an insertion time and a SeMI attack that together achieve the smallest possible α and β.

Threat model. The adversary knows the target z^* and the phase τ in which it may have been inserted. It observes the T released models and cannot influence the training. White-box access gives it the released parameters, and black-box access allows querying the loss of a datapoint at each release. The white-box adversary also knows the training procedure and has reference data, from which it estimates the gradient mean and covariance. The black-box attacks use neither. The auditor is the model owner or a party allowed to insert data, and it runs both the crafter and the adversary. As crafter, it chooses z^* and τ and performs the insertion (Algorithm 1). It then runs one of the attacks above as adversary on each of R rounds and turns the measured error rates into a lower bound on ε (Algorithm 2). The adversary's knowledge of τ is therefore a consequence of the auditor's protocol. Previous auditing works rely on the same canary insertion (Shokri et al., 2017; Nasr et al., 2021; Carlini et al., 2022a; Maddock et al., 2023; Nasr et al., 2023; Steinke et al., 2023), and we introduce the choice of the insertion time. Outside auditing, τ can be known from a subscription date or a federated round. When τ is unknown, SeMI$^{\max}$ and SeMI^{Unif} test whether z^* was inserted at all.

2.2 From Hypothesis Tests to SeMI Attacks

The connection between SeMI attacks and statistical hypothesis testing guides the design of effective adversaries. For the theoretical analysis, we adopt the white-box setting (Sablayrolles et al., 2019; Carlini et al., 2022a; Nasr et al., 2023; Maddock et al., 2023), where the adversary knows the data

[1] For readability we suppress the dependence on z^*, since our focus is on where to insert it.

distributions $(\mathcal{D}_t)$, mechanisms $(\mathcal{M}_t)$, batch size n, and total steps T. Our auditing procedure carries over to other threat models: the form of o_t changes (e.g., query outputs under black-box API access) and the adversary's test adapts accordingly, while the overall procedure remains the same.

As noted earlier, a SeMI attack that knows the insertion time τ corresponds to a statistical test for (2) based on the observation $o_{1:T} = \mathcal{M}(D)$. When the data-generating distributions $\mathcal{D}_1, \ldots, \mathcal{D}_T$ are known, the Neyman–Pearson lemma (Neyman and Pearson, 1933) establishes that a Likelihood Ratio (LR) test is optimal for this simple hypothesis testing problem. Computing the LR associated with that test, under some assumptions on the distributions, is the main approach that we advocate in this paper, leading to SeMI^* for the empirical mean mechanism (Section 3.1) and to SeMI^{SGD} for the DP-SGD mechanism (Section 4).

However, we also explore SeMI attacks that are agnostic to the insertion time but can exploit the sequential nature of data collection. Such attacks $\mathcal{A}(z^*, o_{1:T})$ only depend on z^* and on the output of the mechanism. We first propose to compute a Generalized Likelihood Ratio (GLR) to test $\mathbf{H_0}$ against $\cup_t \mathbf{H}_1^t : (\exists \tau : D \sim \mathcal{P}_1^\tau)$ leading to the $\text{SeMI}^{\max}$ and $\text{SeMI}^{\text{SGD}}_{\max}$ attacks in our two settings. Finally, for the empirical mean mechanism, we also compute $\text{SeMI}^{\text{Unif}}$, which is defined as a Likelihood Ratio test for $\mathbf{H_0}$ against $\widetilde{\mathbf{H}_1} : \ (D \sim \widetilde{\mathcal{P}_1})$ where $\widetilde{\mathcal{P}_1}$ is defined as: sample $\tau \sim \mathcal{U}(\{1, \ldots, T\})$ and then given τ sample D from $\mathcal{P}_1^\tau$.

In the next section, we see that both types of approaches lead to tighter bounds on ε than the optimal MI attack with access to the final model only.

3 SeMI Attacks for the Empirical Mean Mechanism

For the empirical mean mechanism with Gaussian data, the likelihood ratio between $\mathbf{H_0}$ and $\mathbf{H_1^\tau}$ depends on the observed sequence only through the batch mean at the insertion step. This isolation property makes the power of the resulting attack, SeMI^*, independent of τ and of T. The same argument carries over to distribution shifts, multiple insertions, and unknown insertion times.

3.1 Optimal Test: SeMI^*

We analyze the empirical mean mechanism $\mathcal{M}^{\text{emp}}$ with $\mathcal{D}_t = \mathcal{D} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ for all t, mean vector $\boldsymbol{\mu} \in \mathbb{R}^d$, and a positive definite covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$. Each batch D_t contains n i.i.d. samples from $\mathcal{D}$.[2] At step t, $\mathcal{M}^{\text{emp}}$ outputs the cumulative empirical mean $\hat{\boldsymbol{\mu}}_t \triangleq \frac{1}{nt} \sum_{j=1}^{t} \sum_{k=1}^{n} \mathbf{X}_{j,k}$, satisfying the recursion $\hat{\boldsymbol{\mu}}_t = (1 - \frac{1}{t})\hat{\boldsymbol{\mu}}_{t-1} + \frac{1}{t}\bar{\mathbf{X}}_t$ with the t-th batch mean $\bar{\mathbf{X}}_t \triangleq \frac{1}{n} \sum_{k=1}^{n} \mathbf{X}_{t,k}$.

Theorem 3.1 (SeMI^* Score: Multivariate Log Likelihood Ratio (LR)). *Let $\mathbf{N}_\tau \triangleq \bar{\mathbf{X}}_\tau - \boldsymbol{\mu}$, and $c(n) \triangleq \frac{n}{n-1}$. The log-LR for testing $\mathbf{H_0}$ against $\mathbf{H_1^\tau}$ is*

$$\log \text{LR}_\tau = \frac{d}{2} \log c(n) - \frac{c(n)}{2} \mathbf{N}_\tau{}^\top \boldsymbol{\Sigma}^{-1} \mathbf{N}_\tau + c(n) \mathbf{N}_\tau{}^\top \boldsymbol{\Sigma}^{-1} (\mathbf{z}^* - \boldsymbol{\mu}) - \frac{m^*}{2(n-1)}, \tag{4}$$

where $m^* = \|\mathbf{z}^* - \boldsymbol{\mu}\|^2_{\boldsymbol{\Sigma}^{-1}}$ is the squared Mahalanobis distance of the target from the population mean.

The test statistic depends on the target only through its Mahalanobis distance m^*. Targets farther from the mean in Mahalanobis distance are more detectable. SeMI^* recovers the asymptotic analysis of Azize and Basu (2025) in the regime $d, n \to \infty$ with $d/n = \mathcal{O}(1)$. Our finite-sample expression (4) adds a quadratic term in $\mathbf{N}_\tau$ and a normalizer term $(\log c(n))$ that vanish as $n \to \infty$. We call SeMI^* the attack that thresholds (4). Appendix E gives the full derivation and error analysis.

Isolation Property and Power of SeMI^*. The full LR over all the observations factorizes as a product of conditional densities over all transitions (Lemma 3.2). For $t \neq \tau$, the conditional distribution of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ is identical under both hypotheses: the target either has not yet been inserted $(t < \tau)$ or its effect is absorbed into $\hat{\mu}_{t-1}$ $(t > \tau)$. Those terms cancel, leaving the transition at τ.

Lemma 3.2 (Isolation Property). *The LR for testing $\mathbf{H_0}$ against $\mathbf{H_1^\tau}$ based on $(\hat{\mu}_1, \ldots, \hat{\mu}_T)$ satisfies*

$$\frac{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_1^\tau})}{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_0})} = \frac{p(\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_1^\tau})}{p(\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_0})}.$$

[2] Though we fix the batch size for ease of exposition, SeMI tackles varying batch sizes.

Specifically, for $d = 1$ with $\mathcal{D} = \mathcal{N}(\mu, \sigma^2)$, the conditional distributions at time τ become

$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_0} \sim \mathcal{N}\left(\tfrac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \tfrac{\mu}{\tau}, \tfrac{\sigma^2}{\tau^2 n}\right),$$
$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}\left(\tfrac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \tfrac{(n-1)\mu+z^*}{\tau n}, \tfrac{\sigma^2}{\tau^2 n c(n)}\right).$$

Observations. 1. *Mean shift.* Under $\mathbf{H_1^\tau}$ the mean shifts by $(z^* - \mu)/(\tau n)$ and the variance drops by $(n-1)/n$, since one sample is fixed. 2. *Isolation.* Given consecutive observations $(\hat{\mu}_{\tau-1}, \hat{\mu}_\tau)$, the batch mean is recovered exactly as $\bar{X}_\tau = \tau\hat{\mu}_\tau - (\tau-1)\hat{\mu}_{\tau-1}$. The likelihood ratio depends on the data only through $\bar{X}_\tau$, independent of the batches before and after and of T itself.

We further analytically present the Type I and Type II errors of SeMI* for $d = 1$ (Appendix D).

Lemma 3.3 (SeMI*'s Type I and Type II Errors)**.** *Define* $\gamma_{\max} = \frac{1}{2}\left[\frac{(z^*-\mu)^2}{\sigma^2} + \log(c(n))\right]$. *(a) For* $\gamma > \gamma_{\max}$, $\alpha(\gamma) = 0$. *(b) For* $\gamma \leqslant \gamma_{\max}$, *we get* $\alpha(\gamma) = \Phi(a + b(\gamma)) - \Phi(a - b(\gamma))$ *and* $\beta(\gamma) = \Phi\Big(\frac{a}{\sqrt{c(n)}} - b(\gamma)\sqrt{c(n)}\Big) + \Phi\Big(-\frac{a}{\sqrt{c(n)}} - b(\gamma)\sqrt{c(n)}\Big)$, *where* $m^* = (z^* - \mu)^2/\sigma^2$, $a = \sqrt{m^* n}$, *and* $b(\gamma) = \sqrt{(n-1)\,(m^* + \log c(n) - 2\gamma)}$.

Both errors depend on n and the signal strength m^* but neither on τ nor on T. This τ-invariance follows from the isolation property: the test operates on the recovered batch mean $\bar{X}_\tau$, whose distribution does not depend on when insertion occurred.

Cost of Observing the Final Release vs. the Sequential Releases. A classic adversary sees only $\hat{\mu}_T$ and tests the membership hypotheses from a single observation. Here, the LR takes the same quadratic form as SeMI* with n replaced by nT because the single release pools nT samples, among which one is the target. Thus, the per-sample contribution shrinks by a factor of T. We call this baseline the *Final Observation* test, and denote its errors $\alpha_{\mathrm{FO}}, \beta_{\mathrm{FO}}$.

Lemma 3.4 (Final Observation MI's Type I and Type II Errors)**.** *With* $\gamma_{\max}^{(T)} = \frac{1}{2}\left[m^* + \log(c(nT))\right]$ *and* $\gamma \leqslant \gamma_{\max}^{(T)}$, $\alpha_{\mathrm{FO}}(\gamma) = \Phi(a_T + b_T(\gamma)) - \Phi(a_T - b_T(\gamma))$, *and* $\beta_{\mathrm{FO}}(\gamma) = \Phi\left(a_T/\sqrt{c(nT)} - b_T(\gamma)\sqrt{c(nT)}\right) + \Phi\left(-a_T/\sqrt{c(nT)} - b_T(\gamma)\sqrt{c(nT)}\right)$, *with* $a_T = \sqrt{m^* nT}$ *and* $b_T(\gamma) = \sqrt{(nT-1)\,(m^* + \log(c(nT)) - 2\gamma)}$.

The target influence $(z^* - \mu)/(nT)$ vanishes as T grows, so the power of Final Observation degrades with the number of training steps. SeMI* operates on a recovered batch of size n and keeps its power constant in T: sequential access is what prevents the signal from diluting.

3.2 Relaxing the Assumptions of SeMI*

Each weakening of the known-τ, stationary, and single-insertion setting leads to the same isolation argument applied to a modified transition. Table 1 summarizes the resulting tests.

Unknown τ. We propose two adversaries that do not know τ but are still grounded in hypothesis testing but consider different alternatives: $\bigcup_t \mathbf{H_1^t}$ for SeMI$^{\max}$ and $\widetilde{\mathbf{H_1}}$ for SeMI^{Unif}. Details on their computations are given in Appendices G and H.

Distribution Shift. When $\mathcal{D}_t = \mathcal{N}(\mu_t, \sigma_t^2)$ varies with t, the same cancellation argument applies: for $t \neq \tau$ the conditional of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ is identical under both hypotheses regardless of (μ_t, σ_t^2). The log-LR depends only on $(\mu_\tau, \sigma_\tau^2)$ and the power is unchanged when the auditor knows these at the insertion step (Appendix F). Intuitively, isolation makes the test blind to every distribution but the one at τ, so the effect of shifts at other steps cancels.

Multiple Insertion Times. The same cancellation extends to a known set $S \subseteq [T]$ of insertion steps, with the alternative $\mathbf{H_1^\tau}$ replaced by a multi-insertion alternative $\mathbf{H_1^S}$. We defer the formal definition and the resulting factorization $\log \mathrm{LR}_S = \sum_{t \in S} \log \mathrm{LR}_t$ to Appendix D.5.

3.3 Numerical Illustration: SeMI versus Final Observation MI

We sample a base dataset of size $N = 20$ from $\mathcal{N}(0, 1)$ and form $T = 10$ batches of size $n = 10$ by subsampling with replacement from this base dataset, so consecutive batches can share entries and a single batch can contain duplicates. The target z^* (Mahalanobis distance 5) is inserted into

Table 1: SeMI variants. All attacks threshold a log-LR statistic on the observed snapshot sequence and rely on the isolation property of Lemma 3.2. LR_t is the single-time likelihood ratio of Theorem 3.1. The SGD variant is introduced in Section 4. (NP = Neyman–Pearson)

Variant	Knowledge of τ	Test statistic	Optimality
SeMI^{*}	known τ	$\log \mathrm{LR}_\tau$	NP-optimal against $\mathbf{H_1^\tau}$
$\mathrm{SeMI}^{\mathrm{Unif}}$	$\tau \sim \mathcal{U}([T])$	$\log \frac{1}{T}\sum_{t=1}^{T} \mathrm{LR}_t$	NP-optimal against $\widetilde{\mathbf{H_1}}$
$\mathrm{SeMI}^{\max}$	unknown	$\max_t \log \mathrm{LR}_t$	GLR against $\bigcup_t \mathbf{H_1^t}$
$\mathrm{SeMI}^{\mathrm{SGD}}$	known τ, (DP-)SGD	Thm. 4.1	NP-optimal for Gaussian-gradient approx.

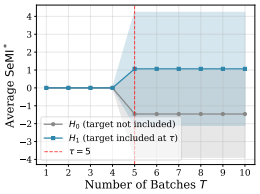

(a) SeMI^{*} statistic

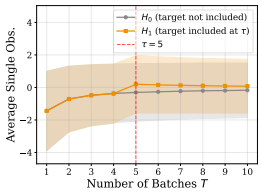

(b) Final Obs. statistic

Figure 2: Log-likelihood ratio as a function of observed snapshots ($n = 10$, $\tau = 5$). (a) SeMI^{*}: separation appears at τ and remains constant. (b) Final Observation: separation appears at τ and decreases as T grows.

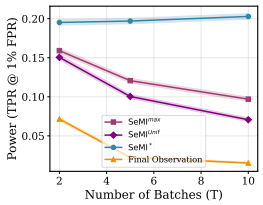

(a) Power vs T

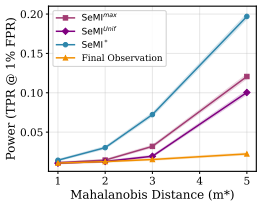

(b) Power vs z^* strength

Figure 3: Comparison of sequential attacks ($n = 10$). (a) Power (TPR at 1% FPR) vs number of snapshots T for z^* with Mahalanobis distance 5. (b) Power vs target Mahalanobis distance as strength of chosen z^* ($T = 5$).

batch $\tau = 5$ only. This subsampled-with-replacement protocol departs from the i.i.d. disjoint-batch assumption underlying SeMI^{*}. The figures below show that the sequential advantage carries through. Errors are estimated from $R = 50\,000$ independent runs.

1. Figure 2 illustrates the *isolation* property: the SeMI^{*} statistic exhibits a sharp step at τ that remains roughly constant in T, while the Final Observation statistic peaks around τ and decays as T grows.

2. Figure 3 quantifies the *power gap*. SeMI^{*} keeps near-constant power across T, while $\mathrm{SeMI}^{\mathrm{Unif}}$ and $\mathrm{SeMI}^{\max}$ degrade with T but stay above Final Observation. The gap widens with the target's Mahalanobis distance.

The matching plots under the disjoint-batch protocol that exactly satisfies the test assumptions, ROC curves, and the multivariate generalization are reported in Appendices J.1, J.2, and J.4.

4 White- and Black-box SeMI Attacks for Gradient Descent

The empirical mean analysis establishes the sequential advantage in an idealized setting. To audit real ML models, we need attacks that apply to gradient descent. Prior sequential attacks on SGD rely on heuristic score combinations (Jagielski et al., 2023). We design a principled LR-based white-box attack and a loss-based black-box attack for SGD that exploit the isolation property observed in the empirical mean case.

I. White-box Attack. In the white-box setting, we observe model parameters $\theta_0, \theta_1, \ldots, \theta_T$ with $\theta_t \in \mathbb{R}^d$. At each step t, $\theta_t = \theta_{t-1} - \eta_t g_t$, where $\eta_t > 0$ is the learning rate and $g_t = \frac{1}{n_t}\sum_{x \in D_t} \nabla_\theta \ell(\theta_{t-1}; x)$ is the batch gradient on D_t of size n_t. As in Section 3, we test $\mathbf{H_0}$ (target z^* not in the training data) against $\mathbf{H_1^\tau}$ (one sample in D_τ replaced by z^*).

We assume that, conditioned on θ_{t-1}, the batch gradient g_t is Gaussian with mean $\mu_g(\theta_{t-1})$ and covariance $\Sigma_g(\theta_{t-1})/n_t$. This approximation can be justified by the central limit theorem when the batch aggregates gradients from many data points. A poor Gaussian fit costs power but leaves the audit valid, since Lemma 2.1 bounds the errors of any test and the lower bound on ε comes from Clopper–Pearson intervals on measured error rates. Gaussian modeling of attack scores is standard (Carlini et al., 2022a; Ye et al., 2022; Leemann et al., 2023), and Appendix J.8 tests it on the batch-gradient means of our experiments. Under this assumption, we observe that the conditional distributions mirror the empirical mean case. Specifically, $\mathbf{H_0} : \theta_t \mid \theta_{t-1} \sim \mathcal{N}\left(\theta_{t-1} - \eta_t \mu_g(\theta_{t-1}), \frac{\eta_t^2 \Sigma_g(\theta_{t-1})}{n_t}\right)$

for all $t \in [T]$, and $\mathbf{H}_1^\tau : \theta_\tau \mid \theta_{\tau-1} \sim \mathcal{N}\left(\theta_{\tau-1} - \eta_\tau \frac{(n_\tau-1)\mu_g + g^*}{n_\tau}, \frac{\eta_\tau^2 (n_\tau - 1)\Sigma_g}{n_\tau^2}\right)$. Here, $g^* \triangleq \nabla_\theta \ell(\theta_{\tau-1}; z^*)$ is the target gradient, and we write μ_g and Σ_g for $\mu_g(\theta_{\tau-1})$ and $\Sigma_g(\theta_{\tau-1})$. The likelihood ratio further factorizes as in Lemma 3.2, where only the transition at time τ contributes.

Theorem 4.1 (SeMI^{SGD} score: Log Likelihood Ratio of Gradients)**.** *Define* $\delta_g = \nabla_\theta \ell(\theta_{\tau-1}; z^*) - \mu_g(\theta_{\tau-1})$ *and* $N = \theta_\tau - \theta_{\tau-1} + \eta_\tau \mu_g(\theta_{\tau-1})$. *The log likelihood ratio is:*

$$\log \text{LR} = -\frac{d}{2} \log\left(\frac{n_\tau - 1}{n_\tau}\right) - \frac{n_\tau}{2(n_\tau - 1)\eta_\tau^2} N^\top \Sigma_g^{-1} N - \frac{n_\tau}{(n_\tau - 1)\eta_\tau} N^\top \Sigma_g^{-1} \delta_g - \frac{m^*}{2(n_\tau - 1)},$$

where $m^* \triangleq \delta_g^\top \Sigma_g^{-1} \delta_g$ *is the squared Mahalanobis distance of the target gradient from the population mean.*

Consequences. *1. Mahalanobis Scores:* The test statistic SeMI^{SGD} has the same structure as Theorem 3.1, with the batch mean replaced by the scaled parameter update N/η_τ. δ_g measures how much the target's gradient deviates from the population mean of the gradient: targets whose gradients differ substantially from typical gradients are more detectable. The Mahalanobis distance m^* accounts for the gradient covariance structure, giving less weight to directions with high variance.

2. *Isolation Property.* From consecutive parameter observations $(\theta_{\tau-1}, \theta_\tau)$, the scaled update N can be computed exactly. The likelihood ratio depends only on this quantity, not on earlier or later parameters. This mirrors the isolation property for the empirical mean. Isolation requires the target to touch exactly one transition and the update to depend only on the previous observed state. It fails if the canary persists into later phases (Appendix D.5), if the auditor cannot see $\theta_{\tau-1}$, or if the optimizer aggregates several updates or the gradient history.

3. *Extension to DP-SGD (Abadi et al., 2016).* For differentially private training via DP-SGD, each update includes clipped gradients and Gaussian noise: $\theta_t = \theta_{t-1} - \eta_t(g_t^{\text{clip}} + \xi_t)$ where $\xi_t \sim \mathcal{N}(0, \sigma_{\text{DP}}^2 C^2 I / n_t^2)$ with clipping threshold C. The noise reduces attack power but preserves the sequential structure and the isolation property. We extend SeMI^{SGD} to DP-SGD by estimating and plugging in three quantities: the population mean gradient $\mu_g(\theta_{\tau-1})$, the gradient covariance $\Sigma_g(\theta_{\tau-1})$, and the target gradient $\nabla_\theta \ell(\theta_{\tau-1}; z^*)$. In our experiments, $\hat{\mu}_g$ and $\hat{\Sigma}_g$ are the sample mean and covariance of the unclipped gradients of 1000 reference samples at initialization, and the target gradient is computed at $\theta_{\tau-1}$. Clipping and DP noise therefore do not enter $\hat{\Sigma}_g$. For larger networks, we estimate the covariance on the last layer only (Appendix J.5), so its cost scales with that layer and not with the network.

II. Black-box Attacks. In this setting, the auditor queries each released checkpoint and observes only the target loss $\ell(\theta_t; z^*)$. We introduce two per-transition scores at the insertion time:

$$\Delta_\tau(z^*) \triangleq \ell(\theta_{\tau-1}; z^*) - \ell(\theta_\tau; z^*), \qquad R_\tau(z^*) \triangleq \frac{\ell(\theta_{\tau-1}; z^*)}{\ell(\theta_\tau; z^*)}. \tag{5}$$

The isolation property of Section 3 concentrates the membership signal at the transition $\theta_{\tau-1} \to \theta_\tau$, so a loss-based statistic should contrast that pair. These scores are inspired by Jagielski et al. (2023): Back-Front contrasts the losses of the first and last models, and Delta thresholds the maximum loss change across consecutive steps. Restricting Back-Front to $(\theta_{\tau-1}, \theta_\tau)$, or evaluating Delta at the inserted step, recovers (5). We call the resulting attacks Loss Diff and Loss Ratio. Taking the maximum of (5) over $\tau \in [T]$ recovers the Delta baselines of Jagielski et al. (2023) for the unknown-τ case.

Numerical Results. Figure 4 reports ROC curves at $T = 5$ for the attacks introduced above and the baselines of Section 5. We observe that *without DP noise,* SeMI^{SGD} *and its max-over-*τ *variants display the best power.* Under DP-SGD at $\varepsilon = 8$, *the noise compresses every curve toward the diagonal but* SeMI^{SGD} *retains a margin over Loss Diff and Loss Ratio* in the low false-positive-rate regime that drives audit tightness. An alternative pretrain-then-finetune setup is reported in Appendix J.5.

5 Sequential Privacy Audits with SeMI

The empirical mean-based experiments in Section 3 establish the sequential advantage of SeMI^* in an idealized setting. Now, we test whether the same advantage transfers to auditing ML models trained with DP-SGD (Abadi et al., 2016). We organize the experiments around four questions.

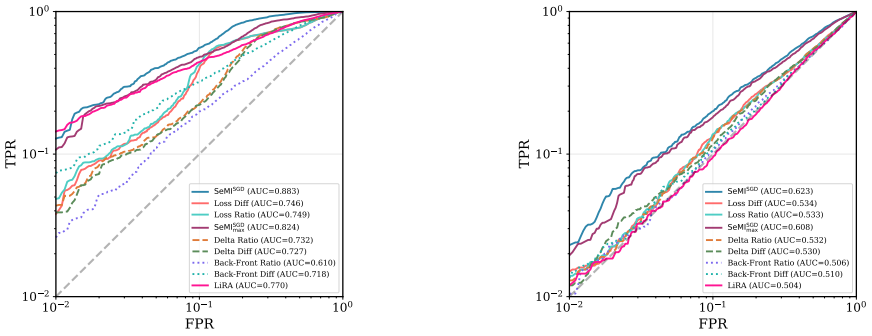

(a) Non-private ($\varepsilon = \infty$) (b) Private ($\varepsilon = 8$)

Figure 4: ROC curves at $T = 5$ snapshots on Fashion-MNIST of SeMI variants and attack baselines.

(Q1) Does observing the full sequence of released snapshots yield tighter ε lower bounds than observing only the final model, and how does the gain scale with the number of releases T? **(Q2)** At fixed T, does the LR-based $\mathrm{SeMI}^{\mathrm{SGD}}$ outperform heuristic scores built from losses or logits? **(Q3)** How much tightness is lost when the insertion time τ is unknown? **(Q4)** What is the impact of the insertion time τ on the audit tightness?

Experimental Setup. A model is trained for L total DP-SGD steps split into T phases of L/T steps, releasing θ_t at every phase boundary. The auditor observes $\theta_0, \ldots, \theta_T$. We keep the single-step statistic of Theorem 4.1 and absorb the L/T steps of a phase into an effective learning rate $\eta_{\text{eff}} = (L/T)\, q\, \eta$, with $q = B/|D|$ and D the full training pool. The dataset is split into T disjoint subsets and each phase is an independent DP-SGD run: weights carry over while the optimizer state and privacy accountant reset, so by parallel composition $\varepsilon_{\text{phase}} = \varepsilon$. Each round inserts z^* in a single fixed phase τ and not in any other. We set $\tau = \lfloor T/2 \rfloor + 1$ for the snapshot-count sweep of Observation 1, and $\tau = 5$ for all other experiments at $T = 5$ in Sections 4 and 5; Observation 4 sweeps τ to isolate its effect.

We consider $T \in \{2, 5, 10\}$, total budget $\varepsilon = 8$, $\delta = 10^{-4}$, clipping norm $C = 1$, and learning rate $\eta = 0.01$. In the main paper, we report softmax regression on Fashion-MNIST ($L = 2500$, $B = 4096$). A fine-tuning variant (pretrain on a public split, then $T = 10$ DP-SGD updates on a private split) on Fashion-MNIST, CIFAR-10, and Purchase-100 is reported in Appendix J.5. Following Azize and Basu (2025), we select an in-distribution canary from the data. Among 1000 candidates drawn from the full dataset, we keep the one whose gradient lies farthest from the data distribution in Mahalanobis distance m^*, estimated on a reference set disjoint from the candidate pool. Theorems 3.1 and 4.1 make the statistic depend on the target only through m^*, so under the Gaussian models of these theorems this choice maximizes power. Each audit runs $R = 1900$ rounds of the SeMI game (Algorithm 2). We compute the empirical privacy lower bound $\hat{\varepsilon}$ via the calibrated-grid protocol of Appendix C.1: $R_{\text{cal}} = 100$ separate calibration rounds define a grid of $K = 99$ thresholds at the centiles of the test statistic, and we apply Clopper–Pearson upper bounds at level $\xi/(2K)$ at every grid point, taking the maximum lower bound across the grid at overall confidence $1 - \xi = 0.98$.

Attack Baselines. We compare three types of methods. (a) *Informed sequential* audits use the knowledge of the insertion time: $\mathrm{SeMI}^{\mathrm{SGD}}$, and the black-box heuristics *Loss Diff* and *Loss Ratio*. (b) *Uninformed sequential* audits use the full snapshot sequence but are oblivious to τ: *Delta* attacks score the maximum loss change across consecutive updates, $\max_t(\ell(\theta_{t-1}; z^*) - \ell(\theta_t; z^*))$ (Diff) or $\max_t \ell(\theta_{t-1}; z^*)/\ell(\theta_t; z^*)$ (Ratio) (Jagielski et al., 2023). *Back-Front* attacks score the total change from initial to final model, $\ell(\theta_0; z^*) - \ell(\theta_T; z^*)$ (Diff) or its ratio (Jagielski et al., 2023). (c) *Snapshot-independent* audits use only the final release: LiRA (Carlini et al., 2022a).

Observation 1: Sequential Advantage from Released Snapshots. Figure 5a reports $\hat{\varepsilon}$ versus the number of released snapshots T at fixed total budget $\varepsilon = 8$ and fixed training steps L. $\mathrm{SeMI}^{\mathrm{SGD}}$ dominates every baseline across T.

Observation 2: Likelihood Ratio vs. Heuristic Score Functions. Figure 6 shows violin plots of the empirical lower bound $\hat{\varepsilon}$ for every attack at $T = 5$ and total budget $\varepsilon = 8$. Each violin is obtained by subsampling 90% of the $R = 1900$ audit rounds without replacement and recomputing $\hat{\varepsilon}$ on each subsample. Among the informed attacks (those that use the insertion time τ), $\mathrm{SeMI}^{\mathrm{SGD}}$ yields the tightest median bound, and its distribution sits above those of Loss Diff and Loss Ratio. All three

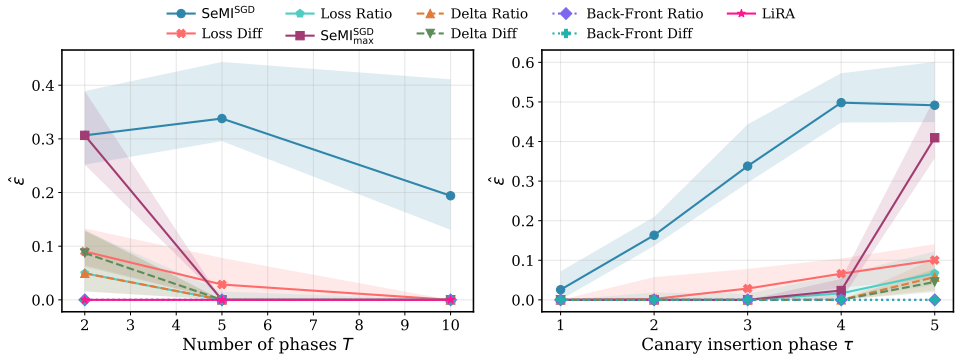

(a) $\hat{\varepsilon}$ vs. #released snapshots T at fixed τ. (b) $\hat{\varepsilon}$ vs. insertion time τ at $T = 5$, when τ is fixed.

Figure 5: Empirical privacy bound $\hat{\varepsilon}$ at $\varepsilon = 8$, $\delta = 10^{-4}$ with fixed training steps L. Higher is tighter.

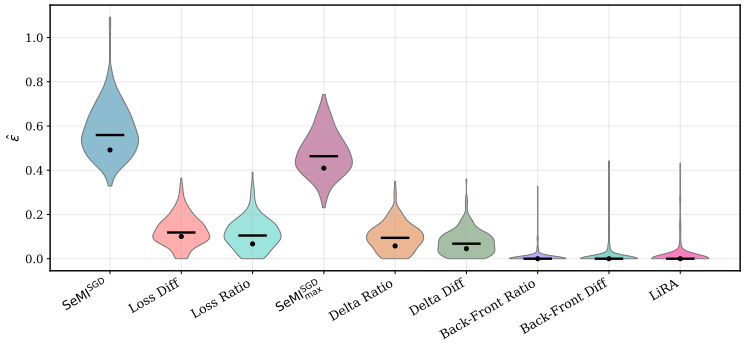

Figure 6: Distribution of empirical epsilon lower bounds at $\delta = 10^{-4}$.

informed attacks contrast $\theta_{\tau-1}$ and θ_τ to use the isolation property of Section 3. The gap reflects that the white-box LR statistic uses the parameter update, which carries more information than the losses at this pair of snapshots. The same ranking holds at $\varepsilon \in \{1, 2, 4\}$ (Appendix J.7).

Observation 3: Knowing the Insertion Time Yields Tighter Audits. Figure 6 pairs each informed attack with its uninformed counterpart: SeMI^{SGD} with $\text{SeMI}^{\text{SGD}}_{\text{max}}$, Loss Diff with Delta Diff, and Loss Ratio with Delta Ratio. In all three pairs the known-τ variant gives a tighter $\hat{\varepsilon}$, mirroring Section 3. The informed attack tests $\mathbf{H_0}$ against $\mathbf{H}_1^\tau$, and when its model holds, no test that ignores τ has higher power against $\mathbf{H}_1^\tau$ at the same level.

Observation 4: Impact of the Insertion Time on the Audit. Sweeping τ across the $T = 5$ phases (Figure 5b) shows that audit tightness depends strongly on τ. For SeMI^{SGD}, later insertion times yield tighter bounds: the per-transition signal is largest once the model has moved away from initialization and the population gradient is small. Choosing a good τ a priori for a given attack and training protocol remains open. We expect it to depend on the learning-rate schedule, which makes some updates more sensitive than others, on how the attack estimates its score, which can tie it to earlier or later updates, and on the canary, which may stop being an outlier later in training.

6 Discussion and Future Work

We derived the optimal MI attack for sequential model observations and characterized its power in closed form under the empirical mean mechanism. The isolation property shows that sequential access recovers the batch statistic at insertion time, removing the dilution that weakens static attacks as training data accumulates. The extension to SGD relies on batch-gradient Gaussianity, yet experiments on DP-SGD training show that SeMI^{SGD} yields tighter privacy lower bounds than heuristic baselines, with the strongest gains in low-privacy regimes.

A promising direction is to adapt single-run audits (Steinke et al., 2023) to the sequential setting by assigning each canary its own insertion phase. Distribution shift is another aspect of successive model releases warranting further study: SeMI^* retains its power when τ is known, but quantifying the effect of an insertion-time distribution chosen by a crafter remains open.

Acknowledgments and Disclosure of Funding

This work has been partially supported by the French National Research Agency (ANR) in the framework of the PEPR IA FOUNDRY project (ANR-23-PEIA-0003). We also acknowledge the Inria-ISI, Kolkata Associate Team "SeRAI", and the ANR JCJC for the REPUBLIC project (ANR-22-CE23-0003-01) for partially supporting the project.

References

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318.

Azize, A. and Basu, D. (2025). Some Targets Are Harder to Identify than Others: Quantifying the Target-dependent Membership Leakage. In Li, Y., Mandt, S., Agrawal, S., and Khan, E., editors, *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pages 397–405. PMLR.

Basu, D. and Chanda, D. (2026). Sublinear algorithms for estimating Wasserstein and TV distances: Applications to fairness and privacy auditing. *Transactions on Machine Learning Research*.

Carlini, N., Chien, S., Nasr, M., Song, S., Terzis, A., and Tramer, F. (2022a). Membership inference attacks from first principles. In *2022 IEEE symposium on security and privacy (SP)*, pages 1897–1914. IEEE.

Carlini, N., Jagielski, M., Zhang, C., Papernot, N., Terzis, A., and Tramer, F. (2022b). The privacy onion effect: Memorization is relative. *Advances in Neural Information Processing Systems*, 35:13263–13276.

Chang, H., Edwards, B., Paul, A. S., and Shokri, R. (2024). Efficient privacy auditing in federated learning. In *33rd USENIX Security Symposium (USENIX Security 24)*, pages 307–323.

Chen, M., Zhang, Z., Wang, T., Backes, M., Humbert, M., and Zhang, Y. (2020). When machine unlearning jeopardizes privacy. *arXiv preprint arXiv:2005.02205*.

Clopper, C. J. and Pearson, E. S. (1934). The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.

Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407.

Homer, N., Szelinger, S., Redman, M., Duggan, D., Tembe, W., Muehling, J., Pearson, J. V., Stephan, D. A., Nelson, S. F., and Craig, D. W. (2008). Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS genetics*, 4(8):e1000167.

Jagielski, M., Ullman, J., and Oprea, A. (2020). Auditing differentially private machine learning: How private is private SGD? *Advances in Neural Information Processing Systems*, 33:22205–22216.

Jagielski, M., Wu, S., Oprea, A., Ullman, J., and Geambasu, R. (2023). How to Combine Membership-Inference Attacks on Multiple Updated Machine Learning Models. *Proceedings on Privacy Enhancing Technologies*, 2023(3):211–232.

Jayaraman, B. and Evans, D. (2019). Evaluating differentially private machine learning in practice. In *28th USENIX Security Symposium*, pages 1895–1912.

Kairouz, P., Oh, S., and Viswanath, P. (2015). The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR.

Kiefer, J. and Wolfowitz, J. (1952). Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, pages 462–466.

Koskela, A. and Mohammadi, J. (2024). Auditing Differential Privacy Guarantees Using Density Estimation.

Krizhevsky, A., Hinton, G., et al. (2009). Learning multiple layers of features from tiny images.

Leemann, T., Pawelczyk, M., and Kasneci, G. (2023). Gaussian membership inference privacy. *Advances in Neural Information Processing Systems*, 36:73866–73878.

Li, H., Li, Z., Wu, S., Hu, C., Ye, Y., Zhang, M., Feng, D., and Zhang, Y. (2024). SeqMIA: Sequential-metric based membership inference attack. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, CCS '24, pages 3496–3510. ACM.

Liu, Y., Zhao, Z., Backes, M., and Zhang, Y. (2022). Membership inference attacks by exploiting loss trajectory. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, pages 2085–2098. ACM.

Maddock, S., Sablayrolles, A., and Stock, P. (2023). CANIFE: Crafting Canaries for Empirical Privacy Measurement in Federated Learning.

Massart, P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *Annals of Probability*, 18.

Nasr, M., Hayes, J., Steinke, T., Balle, B., Tramèr, F., Jagielski, M., Carlini, N., and Terzis, A. (2023). Tight Auditing of Differentially Private Machine Learning.

Nasr, M., Songi, S., Thakurta, A., Papernot, N., and Carlin, N. (2021). Adversary instantiation: Lower bounds for differentially private machine learning. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 866–882. IEEE.

Neyman, J. and Pearson, E. S. (1933). IX. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231(694-706):289–337.

Sablayrolles, A., Douze, M., Schmid, C., Ollivier, Y., and Jégou, H. (2019). White-box vs black-box: Bayes optimal strategies for membership inference. In *International Conference on Machine Learning*, pages 5558–5567. PMLR.

Salem, A., Bhattacharya, A., Backes, M., Fritz, M., and Zhang, Y. (2020). Updates-leak: Data set inference and reconstruction attacks in online learning. In *29th USENIX Security Symposium*, pages 1291–1308.

Shokri, R., Stronati, M., Song, C., and Shmatikov, V. (2017). Membership inference attacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages 3–18. IEEE.

Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Steinke, T., Nasr, M., and Jagielski, M. (2023). Privacy Auditing with One (1) Training Run.

Tramer, F., Terzis, A., Steinke, T., Song, S., Jagielski, M., and Carlini, N. (2022). Debugging differential privacy: A case study for privacy auditing. *arXiv preprint arXiv:2202.12219*.

Wasserman, L. and Zhou, S. (2010). A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105(489):375–389.

Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.

Ye, J., Maddi, A., Murakonda, S. K., Bindschaedler, V., and Shokri, R. (2022). Enhanced membership inference attacks against machine learning models. In *Proceedings of the 2022 ACM SIGSAC conference on computer and communications security*, pages 3093–3106.

Zanella-Béguelin, S., Wutschitz, L., Tople, S., Rühle, V., Paverd, A., Ohrimenko, O., Köpf, B., and Brockschmidt, M. (2020). Analyzing information leakage of updates to natural language models. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*, pages 363–375.

A Connection to Literature

Connection to Jagielski et al. (2023). Appendix B.2 of Jagielski et al. (2023) contains a Neyman–Pearson-optimal distinguisher for membership inference on mean estimation, overlapping in spirit with our derivation of SeMI^*. Their analysis considers two released estimates $(\hat{\mu}_0, \hat{\mu}_1)$ and yields asymptotic concentration bounds; we treat the full sequence of T snapshots, derive the joint likelihood ratio from which the isolation property follows, and obtain exact non-asymptotic expressions for the Type-I and Type-II errors. Their tests ask whether a target was included in any update, or when it was included; the Appendix B.2 distinguisher addresses only inclusion. Because our auditing procedure inserts the target at a chosen time, our primary test SeMI^* asks whether insertion occurred at that specific τ. We also address the inclusion question through $\mathrm{SeMI}^{\mathrm{Unif}}$ (prior on τ) and $\mathrm{SeMI}^{\max}$ (no prior). Their likelihood-ratio analysis is confined to mean estimation; we extend the principled derivation to SGD via a batch-gradient Gaussianity approximation, at the cost of an additional modelling assumption that their distribution-free bounds avoid.

Connection to Chang et al. (2024). Chang et al. (2024) propose the free-training attack (FTA) for membership inference in federated learning, thresholding a closed-form (least-squares slope) aggregation of per-sample performance metrics across FL rounds. The statistic is heuristic and carries no likelihood-ratio optimality guarantee. Intuitively, our isolation property explains why slope aggregation across many rounds dilutes the signal that is in fact concentrated at the insertion round. FTA asks only whether a target was ever included, and its analysis is specific to FL with held-out evaluation; SeMI^*'s isolation argument applies to any sequential mechanism satisfying the Gaussian-batch model, at the price of that modelling assumption.

Connection to Nasr et al. (2023). Nasr et al. (2023) audit DP-SGD by inserting a canary at a chosen batch and showing empirically that observing per-step gradients or model snapshots tightens the resulting ε estimates. Our analysis explains the empirical gain through the isolation property: consecutive snapshots isolate the per-step update at τ (exactly in the empirical-mean case, up to the per-step DP noise under DP-SGD). Intuitively, audit power then scales with the per-step noise rather than the aggregate noise over training (Theorem 4.1). SeMI also covers settings their procedure does not: it does not require per-step gradient access and applies to any sequence of released models, such as public checkpoints or pretrained-then-finetuned releases.

Connection to SeqMIA (Li et al., 2024). SeqMIA and SeMI both exploit a sequence of observations, but the sequences come from different places. SeqMIA is black-box on a single released model. It distills the target, computes several metrics on snapshots of the student, and classifies the resulting sequence with an attention-based RNN trained on shadow models, which decides whether the target belongs to the training set. SeMI reads the sequence of models that the mechanism released and tests whether the target was inserted in a given phase. Among our tests, the counterparts of SeqMIA are the black-box Loss Diff and Loss Ratio. Neither attack contains the other, as SeqMIA never sees a genuine intermediate model and our tests do not apply to a single release.

Dot-product scores. The leading term $N^\top \Sigma_g^{-1} \delta_g$ in Theorem 4.1 is a Σ_g^{-1}-weighted inner product between the scaled parameter update and the target gradient deviation. This parallels gradient-canary auditing methods such as CANIFE (Maddock et al., 2023), which craft a canary gradient and score its inner product with the observed parameter update; the use of the Mahalanobis weighting connects to the asymptotically optimal MI analysis of Azize and Basu (2025). The sequential adaptation restricts the inner product to the per-transition update N between $\theta_{\tau-1}$ and θ_τ, concentrating the score on the snapshot pair where the target's signal sits. $\mathrm{SeMI}^{\mathrm{SGD}}$ derives the precise weighting and target gradient deviation from the Gaussian-gradient model, recovering the heuristic dot-product score as its leading term and giving non-asymptotic guarantees in finite samples rather than only in the limit.

B Proof of the Audit Lemma

This appendix proves Lemma 2.1 in the sequential setting.

Note that fixing τ does not weaken the DP guarantee. By assumption, $\mathcal{M}$ is (ε, δ)-DP with respect to the joint dataset $D = \bigcup_t D_t$: inputs differing in one element yield close outputs, regardless of the batch that element sits in. The crafter's choice of τ isolates a specific neighbouring pair, since under $B = 0$ and $B = 1$ the datasets differ in one element located in batch D_τ. DP holds for this pair, and Lemma 2.1 bounds the errors of *any* test distinguishing the two cases, including tests whose decision rule depends on τ.

Lemma 2.1 (Connecting Sequential MI Attack Error and DP Audits)**.** *Fix a target $z^* \in \mathcal{Z}$ and an insertion time τ. For any sequential MI attack $\mathcal{A}$, if the mechanism $\mathcal{M}$ is (ε, δ)-DP, then*

$$\alpha(\tau, \mathcal{A}) + e^{\varepsilon}\beta(\tau, \mathcal{A}) \geqslant 1 - \delta, \qquad \beta(\tau, \mathcal{A}) + e^{\varepsilon}\alpha(\tau, \mathcal{A}) \geqslant 1 - \delta.$$

Proof. Fix $z^* \in \mathcal{Z}$ and $\tau \in [T]$.

Sample $X_{t,k} \overset{\text{iid}}{\sim} \mathcal{D}_t$ for $t \in [T]$, $k \in [n]$, and $J \sim \mathcal{U}([n])$, jointly independent. Set $D_t = (X_{t,1}, \ldots, X_{t,n})$ and $D = \bigcup_{t=1}^{T} D_t$. Define D' to coincide with D except that the J-th entry of D_τ is replaced by z^*. By construction, $D \sim \mathcal{P}_0$ and $D' \sim \mathcal{P}_1^\tau$, and the two datasets differ in exactly one element.

Let $\Phi \triangleq \mathcal{A}(z^*, \tau, \cdot) \circ \mathcal{M}$ denote the composed (and possibly randomized) mapping from a dataset to $\widehat{B} \in \{0, 1\}$. DP is preserved under post-processing (Dwork and Roth, 2014), so Φ is (ε, δ)-DP whenever $\mathcal{M}$ is.

By construction, for every realization ω, the pair $(D(\omega), D'(\omega))$ differs in exactly one element (the $J(\omega)$-th entry of batch τ), so it is a valid pair of neighboring datasets. Applying the (ε, δ)-DP guarantee of Φ to the event $\{\widehat{B} = 0\}$ at the fixed neighboring pair $(D(\omega), D'(\omega))$ gives, for every ω,

$$\mathbb{P}_\Phi[\Phi(D(\omega)) = 0] \leqslant e^{\varepsilon}\, \mathbb{P}_\Phi[\Phi(D'(\omega)) = 0] + \delta,$$

where the probability is over Φ's internal randomness only. Taking expectation over ω (i.e., over the coupling (D, D', J)) and using $D \sim \mathcal{P}_0$, $D' \sim \mathcal{P}_1^\tau$ marginally,

$$\mathbb{P}_{D \sim \mathcal{P}_0}[\Phi(D) = 0] \leqslant e^{\varepsilon}\, \mathbb{P}_{D' \sim \mathcal{P}_1^\tau}[\Phi(D') = 0] + \delta.$$

By definition, $\mathbb{P}_{\mathcal{P}_0}[\Phi = 0] = 1 - \alpha(\tau, \mathcal{A})$ and $\mathbb{P}_{\mathcal{P}_1^\tau}[\Phi = 0] = \beta(\tau, \mathcal{A})$. Substituting yields

$$1 - \alpha(\tau, \mathcal{A}) \leqslant e^{\varepsilon}\, \beta(\tau, \mathcal{A}) + \delta,$$

which rearranges to $\alpha(\tau, \mathcal{A}) + e^{\varepsilon}\beta(\tau, \mathcal{A}) \geqslant 1 - \delta$.

For the second inequality, apply DP in the reverse direction: the neighboring pair is $(D'(\omega), D(\omega))$ and the event is $\{\widehat{B} = 1\}$. For every ω,

$$\mathbb{P}_\Phi[\Phi(D'(\omega)) = 1] \leqslant e^{\varepsilon}\, \mathbb{P}_\Phi[\Phi(D(\omega)) = 1] + \delta.$$

Taking expectation over ω gives $\mathbb{P}_{\mathcal{P}_1^\tau}[\Phi = 1] \leqslant e^{\varepsilon}\, \mathbb{P}_{\mathcal{P}_0}[\Phi = 1] + \delta$, i.e.,

$$1 - \beta(\tau, \mathcal{A}) \leqslant e^{\varepsilon}\, \alpha(\tau, \mathcal{A}) + \delta,$$

which rearranges to $\beta(\tau, \mathcal{A}) + e^{\varepsilon}\alpha(\tau, \mathcal{A}) \geqslant 1 - \delta$. □

C Audit Protocol and Lower Bound on ε

This appendix details the audit protocol used in our experiments (Section C.1) and a calibration-free DKW alternative (Section C.2). Both produce a high-probability lower bound on ε from rounds of the SeMI game at a fixed insertion time τ, by plugging concentration-based upper confidence bounds on the test errors into (3).

Test family and setup. All tests that we consider have the form $\mathcal{A}^\gamma(z^*, \tau, o_{1:T}) = \mathbb{1}(S > \gamma)$, where $S = s(z^*, \tau, o_{1:T})$ is a real-valued test statistic and γ a threshold. We run R rounds of the SeMI game (Algorithm 1) and record membership indicators $B_1, \dots, B_R$ and statistic values $S_1, \dots, S_R$ rather than the decisions $\widehat{B}_1^\gamma, \dots, \widehat{B}_R^\gamma$ at a fixed γ. Let $N_i(R) \triangleq \sum_{r=1}^R \mathbb{1}(B_r = i)$ count rounds with membership label $i \in \{0, 1\}$. For any $\gamma \in \mathbb{R}$, the empirical Type-I and Type-II errors are

$$\widehat{\alpha}_R(\gamma) = \frac{1}{N_0(R)} \sum_{r=1}^R \mathbb{1}(S_r > \gamma)\, \mathbb{1}(B_r = 0), \qquad \widehat{\beta}_R(\gamma) = \frac{1}{N_1(R)} \sum_{r=1}^R \mathbb{1}(S_r \leqslant \gamma)\, \mathbb{1}(B_r = 1).$$

C.1 Clopper–Pearson with a calibrated threshold grid

This is the protocol used in our experiments. It splits the audit budget into a calibration phase and an evaluation phase, so the thresholds at which Clopper–Pearson intervals are computed are independent of the data used to estimate the errors.

Calibration. We run R_{cal} independent SeMI rounds and collect the resulting statistic values. The threshold grid $\gamma_1 < \cdots < \gamma_K$ is the set of K empirical quantiles of these values at levels $1/(K+1), \dots, K/(K+1)$. The calibration runs are disjoint from the audit runs.

Pointwise upper bounds. For $x \in \{0, \dots, m\}$ successes out of m Bernoulli trials and confidence level $\eta \in (0, 1)$, the Clopper–Pearson upper limit (Clopper and Pearson, 1934) is

$$\mathrm{CP}_{\mathrm{up}}(x, m, \eta) \triangleq \inf\big\{p \in [0, 1] : \mathbb{P}_{Y \sim \mathrm{Binomial}(m,p)}(Y \leqslant x) \leqslant \eta\big\}, \tag{6}$$

with the convention $\mathrm{CP}_{\mathrm{up}}(m, m, \eta) = 1$. Equivalently, $\mathrm{CP}_{\mathrm{up}}(x, m, \eta)$ is the $1 - \eta$ quantile of $\mathrm{Beta}(x + 1,\, m - x)$. By construction, $\mathbb{P}_{Y \sim \mathrm{Binomial}(m,p)}(\mathrm{CP}_{\mathrm{up}}(Y, m, \eta) < p) \leqslant \eta$ for every $p \in [0, 1]$ and every $m \geqslant 1$.

Fix $\xi \in (0, 1)$. Conditional on the calibration set and on $(N_0(R), N_1(R))$, the count $N_0(R)\, \widehat{\alpha}_R(\gamma_k)$ is $\mathrm{Binomial}(N_0(R),\, \alpha(\gamma_k))$ and the count $N_1(R)\, \widehat{\beta}_R(\gamma_k)$ is $\mathrm{Binomial}(N_1(R),\, \beta(\gamma_k))$. We apply Clopper–Pearson upper limits at level $\xi/(2K)$,

$$\overline{\alpha}_R^{\xi}(\gamma_k) = \mathrm{CP}_{\mathrm{up}}\big(N_0(R)\, \widehat{\alpha}_R(\gamma_k),\, N_0(R),\, \xi/(2K)\big),$$
$$\overline{\beta}_R^{\xi}(\gamma_k) = \mathrm{CP}_{\mathrm{up}}\big(N_1(R)\, \widehat{\beta}_R(\gamma_k),\, N_1(R),\, \xi/(2K)\big),$$

with the convention $\overline{\alpha}_R^{\xi}(\gamma_k) = 1$ when $N_0(R) = 0$ and $\overline{\beta}_R^{\xi}(\gamma_k) = 1$ when $N_1(R) = 0$, so that the bound is trivially valid on these zero-probability events. By (6), $\mathbb{P}\big(\alpha(\gamma_k) > \overline{\alpha}_R^{\xi}(\gamma_k) \;\big|\; \text{calibration}\big) \leqslant \xi/(2K)$, and likewise for $\beta(\gamma_k)$.

Audit lower bound. The factor $1/(2K)$ leaves room for a union bound over the K thresholds and the two error types. Plugging the upper bounds into (3) and maximizing over the grid yields

$$\underline{\varepsilon}_R^{\xi,\mathrm{CP}}(\delta) = \log\left(\max_{k \in [K]} \max\left[\frac{1 - \delta - \overline{\alpha}_R^{\xi}(\gamma_k)}{\overline{\beta}_R^{\xi}(\gamma_k)},\; \frac{1 - \delta - \overline{\beta}_R^{\xi}(\gamma_k)}{\overline{\alpha}_R^{\xi}(\gamma_k)}\right]\right). \tag{7}$$

Lemma C.1. *If the mechanism used in SeMI satisfies (ε, δ)-DP, then $\mathbb{P}\big(\varepsilon \geqslant \underline{\varepsilon}_R^{\xi,\mathrm{CP}}(\delta)\big) \geqslant 1 - \xi$, where the probability is jointly over the calibration rounds and the audit rounds.*

Proof. Condition on the calibration set, so that $\gamma_1, \dots, \gamma_K$ are fixed. By the Clopper–Pearson coverage property and a union bound over the $2K$ events,

$$\mathbb{P}\big(E \;\big|\; \text{calibration}\big) \geqslant 1 - \xi, \qquad E \triangleq \big\{\forall\, k \in [K] :\; \alpha(\gamma_k) \leqslant \overline{\alpha}_R^{\xi}(\gamma_k),\; \beta(\gamma_k) \leqslant \overline{\beta}_R^{\xi}(\gamma_k)\big\}.$$

The conditional bound holds for every realization of the calibration set, so it holds marginally: $\mathbb{P}(E) \geqslant 1 - \xi$. On E, fix any $k \in [K]$. Lemma 2.1 gives $\alpha(\gamma_k) + e^{\varepsilon}\beta(\gamma_k) \geqslant 1 - \delta$. If $\beta(\gamma_k) > 0$, dividing through and using $\alpha(\gamma_k) \leqslant \overline{\alpha}_R^{\xi}(\gamma_k)$, $\beta(\gamma_k) \leqslant \overline{\beta}_R^{\xi}(\gamma_k)$ yields $e^{\varepsilon} \geqslant (1 - \delta - \overline{\alpha}_R^{\xi}(\gamma_k))/\overline{\beta}_R^{\xi}(\gamma_k)$. If $\beta(\gamma_k) = 0$, the right-hand side is non-positive on E and the inequality holds vacuously. The symmetric inequality of Lemma 2.1 treats the case with the roles of α and β exchanged. Taking the maximum over k and over the two forms, then applying $\log$, gives $\varepsilon \geqslant \underline{\varepsilon}_R^{\xi,\mathrm{CP}}(\delta)$ on E. □

C.2 DKW envelope: a calibration-free alternative

The protocol above pays a $\log K$ factor in each Clopper–Pearson width and requires a separate calibration budget. An alternative we have not seen used in privacy auditing to our knowledge replaces the pointwise bounds and the union bound over γ with a uniform DKW envelope, which gives a single confidence band valid for every $\gamma \in \mathbb{R}$ and removes the need to calibrate.

Given $\xi > 0$, define DKW upper bounds on the errors

$$\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma) = \widehat{\alpha}_R(\gamma) + \sqrt{\frac{\log(2/\xi)}{2N_0(R)}}, \qquad \overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma) = \widehat{\beta}_R(\gamma) + \sqrt{\frac{\log(2/\xi)}{2N_1(R)}},$$

and the audit lower bound

$$\underline{\varepsilon}_R^{\xi,\mathrm{DKW}}(\delta) = \log\left(\max_{\gamma\in\mathbb{R}} \max\left[\frac{1-\delta-\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)}{\overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)},\ \frac{1-\delta-\overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)}{\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)}\right]\right). \tag{8}$$

Lemma C.2. *If the mechanism used in SeMI satisfies (ε,δ)-DP, then $\mathbb{P}\big(\varepsilon \geqslant \underline{\varepsilon}_R^{\xi,\mathrm{DKW}}(\delta)\big) \geqslant 1-\xi$.*

Proof. Let $F_i(\gamma) \triangleq \mathbb{P}(S_r \leqslant \gamma \mid B_r = i)$ for $i \in \{0,1\}$ be the population CDFs of the test statistic under each hypothesis, and let

$$\hat{F}_i(\gamma) \triangleq \frac{1}{N_i(R)} \sum_{r=1}^{R} \mathbb{1}(S_r \leqslant \gamma)\, \mathbb{1}(B_r = i)$$

be their empirical counterparts on the rounds with $B_r = i$. Then $\alpha(\gamma) = 1-F_0(\gamma)$ and $\beta(\gamma) = F_1(\gamma)$, with $\widehat{\alpha}_R(\gamma) = 1-\hat{F}_0(\gamma)$ and $\widehat{\beta}_R(\gamma) = \hat{F}_1(\gamma)$.

Step 1: uniform concentration. Fix $n_1 \geqslant 1$. Given $N_1(R) = n_1$, the n_1 statistics $\{S_r : B_r = 1\}$ are i.i.d. with CDF F_1, and $\hat{F}_1$ reduces on this conditioning event to the empirical CDF of these n_1 samples, which we denote $\hat{F}_1^{n_1}$. By the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality (Massart, 1990),

$$\mathbb{P}\left(\sup_{\gamma}\big(F_1(\gamma) - \hat{F}_1^{n_1}(\gamma)\big) > t \;\middle|\; N_1(R) = n_1\right) \leqslant e^{-2n_1t^2}.$$

Setting $t = t(n_1) \triangleq \sqrt{\log(2/\xi)/(2n_1)}$ makes the bound equal $\xi/2$. Marginalizing over $N_1(R)$,

$$\begin{aligned}\mathbb{P}\left(\sup_{\gamma}\big(\beta(\gamma) - \widehat{\beta}_R(\gamma)\big) > t(N_1(R))\right) &= \sum_{n_1=1}^{\infty} \mathbb{P}(N_1(R) = n_1) \cdot e^{-2n_1 t(n_1)^2} \\ &= \frac{\xi}{2}\,\mathbb{P}(N_1(R) \geqslant 1) \leqslant \frac{\xi}{2}.\end{aligned}$$

(When $N_1(R) = 0$ we set $\overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma) = 1$. Since $\beta(\gamma) \in [0,1]$, the event $\{\beta(\gamma) > \overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)\}$ is empty, and the resulting lower bound on ε is trivial. The same convention applies to $\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)$ when $N_0(R) = 0$.) Since $\alpha(\gamma) - \widehat{\alpha}_R(\gamma) = \hat{F}_0(\gamma) - F_0(\gamma)$, the same argument applied to the group $\{S_r : B_r = 0\}$ gives $\mathbb{P}(\sup_\gamma(\alpha(\gamma) - \widehat{\alpha}_R(\gamma)) > t(N_0(R))) \leqslant \xi/2$. A union bound over the two groups yields

$$\mathbb{P}(E) \geqslant 1-\xi, \qquad E \triangleq \big\{\forall\,\gamma \in \mathbb{R} :\ \alpha(\gamma) \leqslant \overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma),\ \beta(\gamma) \leqslant \overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)\big\}.$$

Step 2: from event E to the lower bound. On E, fix any $\gamma \in \mathbb{R}$. By Lemma 2.1,

$$\alpha(\gamma) + e^{\varepsilon}\beta(\gamma) \geqslant 1-\delta. \tag{9}$$

We split on the sign of $\beta(\gamma)$. If $\beta(\gamma) > 0$, dividing through gives $e^{\varepsilon} \geqslant (1-\delta-\alpha(\gamma))/\beta(\gamma)$. On E, $\alpha(\gamma) \leqslant \overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)$ and $\beta(\gamma) \leqslant \overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)$, with $\overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma) > 0$ whenever $N_1(R) \geqslant 1$, so

$$e^{\varepsilon} \;\geqslant\; \frac{1-\delta-\alpha(\gamma)}{\beta(\gamma)} \;\geqslant\; \frac{1-\delta-\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)}{\overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma)}. \tag{10}$$

If instead $\beta(\gamma) = 0$, (9) reduces to $\alpha(\gamma) \geqslant 1 - \delta$, and on E the upper bound $\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma) \geqslant \alpha(\gamma) \geqslant 1 - \delta$ makes the right-hand side of (10) non-positive. The inequality then holds vacuously, so (10) is valid for every γ.

The symmetric inequality of Lemma 2.1 yields $e^{\varepsilon} \geqslant (1 - \delta - \overline{\beta}_R^{\xi,\mathrm{DKW}}(\gamma))/\overline{\alpha}_R^{\xi,\mathrm{DKW}}(\gamma)$ by the same argument with the roles of α and β exchanged. Taking the maximum over γ and over the two forms, then applying log, gives $\varepsilon \geqslant \underline{\varepsilon}_R^{\xi,\mathrm{DKW}}(\delta)$ on E. Hence $\mathbb{P}(\varepsilon \geqslant \underline{\varepsilon}_R^{\xi,\mathrm{DKW}}(\delta)) \geqslant \mathbb{P}(E) \geqslant 1 - \xi$. □

D Proofs for SeMI*

We consider T batches of size n drawn from $\mathcal{D} = \mathcal{N}(\mu, \sigma^2)$ with known parameters (μ, σ^2). The target $z^* \in \mathbb{R}$ is fixed. We test $\mathbf{H_0}$ against $\mathbf{H_1^\tau}$, where $\mathbf{H_0}$ is the null hypothesis that z^* is not in D_τ and $\mathbf{H_1^\tau}$ is the alternative hypothesis that z^* is in D_τ. The insertion time $\tau \in \{1, \ldots, T\}$ is assumed known.

Notation. We write $\bar{X}_t = \frac{1}{n}\sum_{k=1}^{n} X_{t,k}$ for the batch mean and $\hat{\mu}_t = \frac{1}{tn}\sum_{j=1}^{t}\sum_{k=1}^{n} X_{j,k}$ for the cumulative empirical mean, which satisfies the recursion $\hat{\mu}_t = \frac{t-1}{t}\hat{\mu}_{t-1} + \frac{1}{t}\bar{X}_t$.

D.1 Proof of Lemma 3.2

Lemma 3.2 (Isolation Property)**.** *The LR for testing* $\mathbf{H_0}$ *against* $\mathbf{H_1^\tau}$ *based on* $(\hat{\mu}_1, \ldots, \hat{\mu}_T)$ *satisfies*

$$\frac{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_1^\tau})}{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_0})} = \frac{p(\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_1^\tau})}{p(\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_0})}.$$

Proof. The likelihood ratio factorizes using the chain rule for conditional densities:

$$\mathrm{LR}_\tau = \frac{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_1^\tau})}{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_0})} = \prod_{t=1}^{T} \frac{p(\hat{\mu}_t \mid \hat{\mu}_1, \ldots, \hat{\mu}_{t-1}, \mathbf{H_1^\tau})}{p(\hat{\mu}_t \mid \hat{\mu}_1, \ldots, \hat{\mu}_{t-1}, \mathbf{H_0})}.$$

The recursive relation $\hat{\mu}_t = \frac{t-1}{t}\hat{\mu}_{t-1} + \frac{1}{t}\bar{X}_t$ implies that, given $\hat{\mu}_{t-1}$, the distribution of $\hat{\mu}_t$ depends only on the distribution of the new batch mean $\bar{X}_t$. Earlier observations $\hat{\mu}_1, \ldots, \hat{\mu}_{t-2}$ provide no additional information once $\hat{\mu}_{t-1}$ is known.

Terms for $t \neq \tau$ cancel. For $t < \tau$, the target z^* has not yet been inserted into any batch. Under both $\mathbf{H_0}$ and $\mathbf{H_1^\tau}$, all samples in batches $D_1, \ldots, D_t$ are drawn i.i.d. from $\mathcal{D}$. The conditional distribution of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ is therefore identical under both hypotheses, and the corresponding factor in the product equals 1.

For $t > \tau$, the batch D_t contains no target sample under either hypothesis. Although the effect of z^* is present in $\hat{\mu}_{t-1}$ under $\mathbf{H_1^\tau}$, conditioning on $\hat{\mu}_{t-1}$ absorbs this effect. The new batch $\bar{X}_t$ is drawn from $\mathcal{N}(\mu, \sigma^2/n)$ under both hypotheses, so the conditional distributions of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ are again identical. These factors also equal 1.

The only factor that differs between the two hypotheses is the one at $t = \tau$, where the batch D_τ contains z^* under $\mathbf{H_1^\tau}$ but not under $\mathbf{H_0}$. □

D.2 Univariate Log Likelihood Ratio

Theorem D.1 (Log Likelihood Ratio, univariate)**.** *The log likelihood ratio for univariate samples has the form*

$$\log \mathrm{LR}_\tau = -\frac{1}{2}\log\left(\frac{n-1}{n}\right) - \frac{n}{2(n-1)\sigma^2}(\bar{X}_\tau - \mu)^2 + \frac{(z^* - \mu)n}{(n-1)\sigma^2}(\bar{X}_\tau - \mu) - \frac{(z^* - \mu)^2}{2(n-1)\sigma^2}.$$

Proof. We derive the log likelihood ratio from the conditional distributions at time τ.

Conditional distributions. From the main text, the conditional distributions are:

$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_0} \sim \mathcal{N}(m_0, v_0)$$
$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}(m_1, v_1)$$

where

$$m_0 = \frac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \frac{\mu}{\tau}, \qquad v_0 = \frac{\sigma^2}{\tau^2 n},$$
$$m_1 = \frac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \frac{(n-1)\mu + z^*}{\tau n}, \qquad v_1 = \frac{(n-1)\sigma^2}{\tau^2 n^2}.$$

Log likelihood ratio. Let $N = \hat{\mu}_\tau - m_0$. The log likelihood ratio between two Gaussians is:

$$\log \mathrm{LR}_\tau = -\frac{1}{2}\log\left(\frac{v_1}{v_0}\right) + \frac{N^2}{2v_0} - \frac{(N-(m_1-m_0))^2}{2v_1}.$$

Expanding $(N-(m_1-m_0))^2$:

$$\log \mathrm{LR}_\tau = -\frac{1}{2}\log\left(\frac{v_1}{v_0}\right) + N^2\left(\frac{1}{2v_0} - \frac{1}{2v_1}\right) + \frac{N(m_1-m_0)}{v_1} - \frac{(m_1-m_0)^2}{2v_1}.$$

Variance ratio and mean difference. The variance ratio is:

$$\frac{v_1}{v_0} = \frac{(n-1)\sigma^2/(\tau^2 n^2)}{\sigma^2/(\tau^2 n)} = \frac{n-1}{n}.$$

The mean difference is:

$$m_1 - m_0 = \frac{(n-1)\mu + z^*}{\tau n} - \frac{\mu}{\tau} = \frac{(n-1)\mu + z^* - n\mu}{\tau n} = \frac{z^*-\mu}{\tau n}.$$

Simplification. Using $N = \frac{1}{\tau}(\bar{X}_\tau - \mu)$ and computing the coefficients:

$$\frac{1}{2v_0} - \frac{1}{2v_1} = \frac{\tau^2 n}{2\sigma^2} - \frac{\tau^2 n^2}{2(n-1)\sigma^2} = -\frac{\tau^2 n}{2(n-1)\sigma^2},$$
$$\frac{m_1-m_0}{v_1} = \frac{(z^*-\mu)\tau n}{(n-1)\sigma^2},$$
$$\frac{(m_1-m_0)^2}{2v_1} = \frac{(z^*-\mu)^2}{2(n-1)\sigma^2}.$$

Substituting $N\tau = \bar{X}_\tau - \mu$, the τ factors cancel in the quadratic and linear terms, yielding:

$$\log \mathrm{LR}_\tau = -\frac{1}{2}\log\left(\frac{n-1}{n}\right) - \frac{n}{2(n-1)\sigma^2}(\bar{X}_\tau-\mu)^2 + \frac{(z^*-\mu)n}{(n-1)\sigma^2}(\bar{X}_\tau-\mu) - \frac{(z^*-\mu)^2}{2(n-1)\sigma^2}.$$

□

D.3 Proof of Lemma 3.3

Lemma 3.3 (SeMI*'s Type I and Type II Errors). *Define* $\gamma_{\max} = \frac{1}{2}\left[\frac{(z^*-\mu)^2}{\sigma^2} + \log(c(n))\right]$. *(a) For* $\gamma > \gamma_{\max}$, $\alpha(\gamma) = 0$. *(b) For* $\gamma \leqslant \gamma_{\max}$, *we get* $\alpha(\gamma) = \Phi(a + b(\gamma)) - \Phi(a - b(\gamma))$ *and* $\beta(\gamma) = \Phi\Big(\frac{a}{\sqrt{c(n)}} - b(\gamma)\sqrt{c(n)}\Big) + \Phi\Big(-\frac{a}{\sqrt{c(n)}} - b(\gamma)\sqrt{c(n)}\Big)$, *where* $m^* = (z^*-\mu)^2/\sigma^2$, $a = \sqrt{m^* n}$, *and* $b(\gamma) = \sqrt{(n-1)\left(m^* + \log c(n) - 2\gamma\right)}$.

Proof. The test statistic has the quadratic form $\log \mathrm{LR}_\tau = c_0 + c_1 N + c_2 N^2$ where $c_2 < 0$.

The condition $\log \mathrm{LR}_\tau \geqslant \gamma$ is equivalent to $c_2 N^2 + c_1 N + (c_0 - \gamma) \geqslant 0$. Since $c_2 < 0$, this downward-opening parabola is non-negative only when N lies between its two roots (if they exist).

The discriminant of the quadratic is:

$$\Delta = c_1^2 - 4c_2(c_0-\gamma) = \frac{\tau^2 n}{\sigma^2(n-1)}\left[\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{n-1}{n}\right) - 2\gamma\right].$$

Case 1: If $\Delta < 0$, the parabola never reaches zero, so $\alpha = 0$. This occurs when $\gamma > \gamma_{\max}$.

Case 2: For $\Delta \geqslant 0$, the roots are:

$$N_{1,2} = \frac{z^* - \mu}{\tau} \mp \frac{\sigma\sqrt{n-1}}{\tau\sqrt{n}} \sqrt{\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{n-1}{n}\right) - 2\gamma}.$$

Under $\mathbf{H_0}$, $N \sim \mathcal{N}(0, \sigma^2/(\tau^2 n))$. The Type I error is:

$$\alpha(\gamma) = P_0(N_1 \leqslant N \leqslant N_2) = \Phi\left(\frac{N_2}{\sigma/(\tau\sqrt{n})}\right) - \Phi\left(\frac{N_1}{\sigma/(\tau\sqrt{n})}\right).$$

Substituting the expressions for N_1 and N_2 and simplifying (noting that τ cancels) yields the stated formula with $a = \frac{(z^*-\mu)\sqrt{n}}{\sigma}$ and $b(\gamma) = \sqrt{n-1}\sqrt{\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{n-1}{n}\right) - 2\gamma}$.

Type II error. The Type II error is $\beta(\gamma) = P_1(\log \mathrm{LR}_\tau < \gamma) = P_1(N < N_1) + P_1(N > N_2)$. Under $\mathbf{H_1^\tau}$, the batch mean is $\bar{X}_\tau \sim \mathcal{N}\left(\frac{(n-1)\mu + z^*}{n}, \frac{(n-1)\sigma^2}{n^2}\right)$, so

$$N = \frac{1}{\tau}(\bar{X}_\tau - \mu) \sim \mathcal{N}\left(\frac{z^* - \mu}{\tau n}, \frac{(n-1)\sigma^2}{\tau^2 n^2}\right).$$

Let $\mu_1 = \frac{z^*-\mu}{\tau n}$ and $\sigma_1 = \frac{\sigma\sqrt{n-1}}{\tau n}$. Then

$$\beta(\gamma) = \Phi\left(\frac{N_1 - \mu_1}{\sigma_1}\right) + \Phi\left(\frac{-(N_2 - \mu_1)}{\sigma_1}\right),$$

and substituting the expressions for $N_{1,2}$ gives

$$\frac{N_{1,2} - \mu_1}{\sigma_1} = \frac{(z^*-\mu)\sqrt{n-1}}{\sigma} \mp \sqrt{n}\sqrt{\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{n-1}{n}\right) - 2\gamma},$$

yielding the stated formula. □

D.4 Proof of Lemma 3.4

Lemma 3.4 (Final Observation MI's Type I and Type II Errors)**.** *With* $\gamma_{\max}^{(T)} = \frac{1}{2}\left[m^* + \log\left(c(nT)\right)\right]$ *and* $\gamma \leqslant \gamma_{\max}^{(T)}$, $\alpha_{\mathrm{FO}}(\gamma) = \Phi(a_T + b_T(\gamma)) - \Phi(a_T - b_T(\gamma))$, *and* $\beta_{\mathrm{FO}}(\gamma) = \Phi\left(a_T/\sqrt{c(nT)} - b_T(\gamma)\sqrt{c(nT)}\right) + \Phi\left(-a_T/\sqrt{c(nT)} - b_T(\gamma)\sqrt{c(nT)}\right)$, *with* $a_T = \sqrt{m^* nT}$ *and* $b_T(\gamma) = \sqrt{(nT-1)\left(m^* + \log(c(nT)) - 2\gamma\right)}$.

Proof. The proof follows the same structure as Lemma 3.3, with n replaced by nT throughout.

For the single-observation case, the test statistic is $\log \mathrm{LR} = c_0^{(T)} + c_1^{(T)} W + c_2^{(T)} W^2$ where $W = \hat{\mu}_T - \mu$ and:

$$c_0^{(T)} = -\frac{1}{2}\log\left(\frac{nT-1}{nT}\right) - \frac{(z^*-\mu)^2}{2(nT-1)\sigma^2},$$
$$c_1^{(T)} = \frac{nT(z^*-\mu)}{(nT-1)\sigma^2},$$
$$c_2^{(T)} = -\frac{nT}{2(nT-1)\sigma^2}.$$

The discriminant is:

$$\Delta = \frac{nT}{\sigma^2(nT-1)}\left[\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{nT-1}{nT}\right) - 2\gamma\right].$$

For $\Delta \geqslant 0$, the roots are:

$$W_{1,2} = \frac{z^*-\mu}{nT} \mp \frac{\sigma\sqrt{nT-1}}{nT}\sqrt{\frac{(z^*-\mu)^2}{\sigma^2} - \log\left(\frac{nT-1}{nT}\right) - 2\gamma}.$$

Under $\mathbf{H_0}$, $W \sim \mathcal{N}(0, \sigma^2/(nT))$, giving the stated Type I error.

Type II error. Under $\mathbf{H_1^\tau}$, one sample among nT is replaced by z^*, so

$$W = \hat{\mu}_T - \mu \sim \mathcal{N}\left(\frac{z^* - \mu}{nT}, \frac{(nT-1)\sigma^2}{(nT)^2}\right).$$

The Type II error $\beta_{\mathrm{FO}}(\gamma) = P_1(W < W_1) + P_1(W > W_2)$ follows by the same calculation as the sequential case in Lemma 3.3, with n replaced by nT. □

D.5 Extension to Multiple Insertion Times

The MI game of Section 2 inserts the target at a single step τ. The cancellation argument of Lemma 3.2 extends to a known set of insertion steps, at the cost of redefining the alternative hypothesis. Let $S \subseteq \{1, \dots, T\}$ be a known set of insertion steps and let $\mathbf{H_1^S}$ denote the hypothesis under which a copy of the target z^* replaces one sample of D_t for each $t \in S$, while each D_t with $t \notin S$ is drawn i.i.d. from $\mathcal{D}$.

Lemma D.2 (Factorization under multiple insertions)**.** *The likelihood ratio for* $\mathbf{H_0}$ *against* $\mathbf{H_1^S}$ *based on* $(\hat{\mu}_1, \dots, \hat{\mu}_T)$ *satisfies*

$$\log \mathrm{LR}_S = \sum_{t \in S} \log \mathrm{LR}_t,$$

where $\log \mathrm{LR}_t$ *is the single-step log-LR from Theorem D.1.*

Proof. The likelihood factorizes as $\prod_{t=1}^{T} p(\hat{\mu}_t \mid \hat{\mu}_{t-1}, H)$. For $t \notin S$, the conditional at step t matches the $t \neq \tau$ argument of Lemma 3.2: either no insertion has occurred by step t, or all past insertions are absorbed into $\hat{\mu}_{t-1}$. Those factors equal 1. Each $t \in S$ contributes the single-step ratio LR_t, and the logarithm is additive. □

E Multivariate Gaussian Extension

We extend the sequential membership inference approach to the multivariate Gaussian setting where datapoints lie in $\mathbb{R}^d$. We consider T batches of size n drawn from $\mathcal{D} = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ where $\boldsymbol{\mu} \in \mathbb{R}^d$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ is positive definite and known. The target $\mathbf{z}^* \in \mathbb{R}^d$ and τ are known. We test $\mathbf{H_0}$ versus $\mathbf{H_1^\tau}$ as defined in (2).

Notation. Bold symbols denote vectors ($\mathbf{x} \in \mathbb{R}^d$) and matrices ($\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$). The Mahalanobis distance is $\|\mathbf{x}\|_{\boldsymbol{\Sigma}} = \sqrt{\mathbf{x}^\top \boldsymbol{\Sigma}^{-1} \mathbf{x}}$.

E.1 Conditional Distributions

The recursive relation $\hat{\boldsymbol{\mu}}_t = \frac{t-1}{t}\hat{\boldsymbol{\mu}}_{t-1} + \frac{1}{t}\bar{\mathbf{X}}_t$ extends naturally.

Lemma E.1 (Multivariate Conditional Distributions)**.** *The conditional distributions at time* τ *are:*

$$\hat{\boldsymbol{\mu}}_\tau \mid \hat{\boldsymbol{\mu}}_{\tau-1}, \mathbf{H_0} \sim \mathcal{N}\left(\frac{\tau-1}{\tau}\hat{\boldsymbol{\mu}}_{\tau-1} + \frac{\boldsymbol{\mu}}{\tau}, \frac{\boldsymbol{\Sigma}}{\tau^2 n}\right),$$

$$\hat{\boldsymbol{\mu}}_\tau \mid \hat{\boldsymbol{\mu}}_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}\left(\frac{\tau-1}{\tau}\hat{\boldsymbol{\mu}}_{\tau-1} + \frac{(n-1)\boldsymbol{\mu} + \mathbf{z}^*}{\tau n}, \frac{(n-1)\boldsymbol{\Sigma}}{\tau^2 n^2}\right).$$

Proof. **Batch mean distributions.** Under $\mathbf{H_0}$, the batch mean is $\bar{\mathbf{X}}_\tau \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}/n)$.

Under $\mathbf{H_1^\tau}$, the batch contains $(n-1)$ samples from $\mathcal{D}$ plus $\mathbf{z}^*$:

$$\bar{\mathbf{X}}_\tau = \frac{1}{n}\left[(n-1)\boldsymbol{\mu}_{\text{sample}} + \mathbf{z}^*\right]$$

where $\boldsymbol{\mu}_{\text{sample}} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}/(n-1))$. The result follows.

The recursion $\hat{\boldsymbol{\mu}}_\tau = \frac{\tau-1}{\tau}\hat{\boldsymbol{\mu}}_{\tau-1} + \frac{1}{\tau}\bar{\mathbf{X}}_\tau$ is an affine transformation, yielding the stated conditional distributions. □

E.2 Log Likelihood Ratio

Theorem E.2 (Multivariate Log Likelihood Ratio)**.** *Define* $\mathbf{N}_\tau = \tau(\hat{\boldsymbol{\mu}}_\tau - \frac{\tau-1}{\tau}\hat{\boldsymbol{\mu}}_{\tau-1} - \frac{\boldsymbol{\mu}}{\tau}) = \bar{\mathbf{X}}_\tau - \boldsymbol{\mu}$. *The log likelihood ratio is:*

$$\log \mathrm{LR}_\tau = -\frac{d}{2}\log\left(\frac{n-1}{n}\right) - \frac{n}{2(n-1)}\mathbf{N}_\tau{}^\top\boldsymbol{\Sigma}^{-1}\mathbf{N}_\tau + \frac{n}{n-1}\mathbf{N}_\tau{}^\top\boldsymbol{\Sigma}^{-1}(\mathbf{z}^* - \boldsymbol{\mu}) - \frac{(\mathbf{z}^* - \boldsymbol{\mu})^\top\boldsymbol{\Sigma}^{-1}(\mathbf{z}^* - \boldsymbol{\mu})}{2(n-1)}.$$

Proof. **Conditional parameters.** Let $\mathbf{m}_0, \mathbf{m}_1$ be the conditional means and $\mathbf{V}_0 = \frac{\boldsymbol{\Sigma}}{\tau^2 n}$, $\mathbf{V}_1 = \frac{(n-1)\boldsymbol{\Sigma}}{\tau^2 n^2}$ the conditional covariances.

Covariance ratio and mean difference. We have $\mathbf{V}_1 = \frac{n-1}{n}\mathbf{V}_0$, so $\det(\mathbf{V}_1) = \left(\frac{n-1}{n}\right)^d \det(\mathbf{V}_0)$ and $\mathbf{V}_1^{-1} = \frac{n}{n-1}\mathbf{V}_0^{-1}$. The mean difference is:

$$\mathbf{m}_1 - \mathbf{m}_0 = \frac{(n-1)\boldsymbol{\mu} + \mathbf{z}^*}{\tau n} - \frac{\boldsymbol{\mu}}{\tau} = \frac{\mathbf{z}^* - \boldsymbol{\mu}}{\tau n}.$$

Log likelihood ratio. The multivariate Gaussian log likelihood ratio is:

$$\log \mathrm{LR}_\tau = -\frac{1}{2}\log\frac{\det(\mathbf{V}_1)}{\det(\mathbf{V}_0)} + \frac{1}{2}\mathbf{L}_\tau{}^\top\mathbf{V}_0^{-1}\mathbf{L}_\tau - \frac{1}{2}(\mathbf{L}_\tau - (\mathbf{m}_1 - \mathbf{m}_0))^\top\mathbf{V}_1^{-1}(\mathbf{L}_\tau - (\mathbf{m}_1 - \mathbf{m}_0))$$

where $\mathbf{L}_\tau = \hat{\boldsymbol{\mu}}_\tau - \mathbf{m}_0 = \frac{1}{\tau}\mathbf{N}_\tau$. The determinant term gives $-\frac{d}{2}\log\left(\frac{n-1}{n}\right)$. Expanding the quadratic terms with $\mathbf{V}_0^{-1} = \tau^2 n\boldsymbol{\Sigma}^{-1}$ and simplifying yields the stated result. □

Corollary E.3 (Quadratic Form)**.** *The log likelihood ratio has the form* $\log \mathrm{LR}_\tau = c_0 + \mathbf{c}_1^\top\mathbf{N}_\tau + \mathbf{N}_\tau{}^\top\mathbf{C}_2\mathbf{N}_\tau$ *where:*

$$c_0 = -\frac{d}{2}\log\left(\frac{n-1}{n}\right) - \frac{m^*}{2(n-1)},$$
$$\mathbf{c}_1 = \frac{n}{n-1}\boldsymbol{\Sigma}^{-1}(\mathbf{z}^* - \boldsymbol{\mu}),$$
$$\mathbf{C}_2 = -\frac{n}{2(n-1)}\boldsymbol{\Sigma}^{-1},$$

and $m^* = (\mathbf{z}^* - \boldsymbol{\mu})^\top\boldsymbol{\Sigma}^{-1}(\mathbf{z}^* - \boldsymbol{\mu})$ *is the squared Mahalanobis distance.*

E.3 Type I and II Error Analysis

Lemma E.4 (Distributions of the Test Statistic)**.** *Under* $\mathbf{H_0}$*:* $\mathbf{N}_\tau \sim \mathcal{N}\left(\mathbf{0}, \frac{\boldsymbol{\Sigma}}{n}\right)$.

Under $\mathbf{H_1^\tau}$*:* $\mathbf{N}_\tau \sim \mathcal{N}\left(\frac{\mathbf{z}^* - \boldsymbol{\mu}}{n}, \frac{(n-1)\boldsymbol{\Sigma}}{n^2}\right)$.

Lemma E.5 (Type I Error Analysis - Diagonal Case)**.** *For the log likelihood ratio test with threshold* γ *in the multivariate Gaussian case with diagonal covariance* $\boldsymbol{\Sigma} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_d^2)$*, let:*

$$\gamma_{\max} = \frac{1}{2}\left[m^* - d\log\left(\frac{n-1}{n}\right)\right]$$

where $m^* = \sum_{j=1}^d \frac{(z_j^* - \mu_j)^2}{\sigma_j^2}$ *is the squared Mahalanobis distance. For* $\gamma \leqslant \gamma_{\max}$*, the Type I error is:*

$$\alpha(\gamma) = \mathbb{P}_0\left(\|\mathbf{U}\|_{\boldsymbol{\Sigma}^{-1}}^2 \leqslant R^2(\gamma)\right)$$

where $\mathbf{U} \sim \mathcal{N}\left(-(\mathbf{z}^* - \boldsymbol{\mu}), \frac{\boldsymbol{\Sigma}}{n}\right)$ *and* $R^2(\gamma) = \frac{2(n-1)}{n}\left[\frac{n^2 m^*}{2} - \gamma - \frac{d}{2}\log\left(\frac{n-1}{n}\right)\right]$.

Proof. With diagonal covariance, the log likelihood ratio sums over dimensions: $\log \mathrm{LR}_\tau = \sum_{j=1}^{d} \log \mathrm{LR}_{\tau,j}$. The sum can be written as $\log \mathrm{LR}_\tau = \mathbf{N}_\tau{}^\top \mathbf{C}_2 \mathbf{N}_\tau + \mathbf{c}_1^\top \mathbf{N}_\tau + c_0$ with diagonal $\mathbf{C}_2$.

Substituting the coefficients from Corollary E.3 (adapted for diagonal $\mathbf{\Sigma}$), we complete the square. Let $\mathbf{M}_\tau = \mathbf{N}_\tau - (\mathbf{z}^* - \boldsymbol{\mu})$. Under $\mathbf{H_0}$, $\mathbf{N}_\tau \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}/n)$, so $\mathbf{M}_\tau \sim \mathcal{N}(-(\mathbf{z}^* - \boldsymbol{\mu}), \mathbf{\Sigma}/n)$.

The condition $\log \mathrm{LR}_\tau \geqslant \gamma$ transforms to:

$$\mathbf{M}_\tau{}^\top \mathbf{C}_2 \mathbf{M}_\tau \geqslant \gamma - c_0 + \frac{1}{4} \mathbf{c}_1^\top \mathbf{C}_2^{-1} \mathbf{c}_1.$$

Substituting the specific values for $c_0, \mathbf{c}_1, \mathbf{C}_2$, the RHS simplifies to $\gamma + \frac{d}{2} \log\left(\frac{n-1}{n}\right) - \frac{n^2 m^*}{2}$.

Since $\mathbf{C}_2 = -\frac{n}{2(n-1)} \mathbf{\Sigma}^{-1}$ is negative definite, the inequality reverses:

$$\mathbf{M}_\tau{}^\top \mathbf{\Sigma}^{-1} \mathbf{M}_\tau \leqslant -\frac{2(n-1)}{n} \left[\gamma + \frac{d}{2} \log\left(\frac{n-1}{n}\right) - \frac{n^2 m^*}{2} \right].$$

Letting $\mathbf{U} = \mathbf{M}_\tau$, the result follows. The probability is the mass of a shifted Gaussian within an ellipsoid defined by the Mahalanobis distance. □

Corollary E.6 (Explicit Type I Error)**.** *For diagonal covariance $\mathbf{\Sigma} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_d^2)$, the Type I error is:*

$$\alpha(\gamma) = F_{\chi^2_d(nm^*)}\left(n\, R^2(\gamma)\right)$$

where $F_{\chi^2_d(\lambda)}$ is the CDF of the non-central chi-squared distribution with d degrees of freedom and non-centrality parameter $\lambda = nm^*$, and $R^2(\gamma)$ is defined in Lemma E.5.

Proof. For diagonal $\mathbf{\Sigma}$, the squared Mahalanobis norm decomposes as $\|\mathbf{U}\|^2_{\mathbf{\Sigma}^{-1}} = \sum_{j=1}^{d} U_j^2/\sigma_j^2$. Each component $U_j \sim \mathcal{N}(-(z_j^* - \mu_j), \sigma_j^2/n)$, so the standardized variable $W_j = \sqrt{n} U_j/\sigma_j \sim \mathcal{N}(-\sqrt{n}(z_j^* - \mu_j)/\sigma_j, 1)$.

The sum $\sum_{j=1}^{d} W_j^2 = n\|\mathbf{U}\|^2_{\mathbf{\Sigma}^{-1}}$ follows a non-central chi-squared distribution with d degrees of freedom and non-centrality parameter $\lambda = \sum_{j=1}^{d} n(z_j^* - \mu_j)^2/\sigma_j^2 = nm^*$. □

F Robustness to Distribution Shift

We consider the setting where each batch may be drawn from a different distribution: $D_t \sim \mathcal{N}(\mu_t, \sigma_t^2)$ for $t = 1, \ldots, T$. The target $z^* \in \mathbb{R}$ and the insertion time τ are fixed. We test $\mathbf{H_0}$ versus $\mathbf{H_1^\tau}$.

Notation. Parameters (μ_t, σ_t^2) may vary with t, but only $(\mu_\tau, \sigma_\tau^2)$ at the insertion time affect the test.

F.1 Main Result

Theorem F.1 (Robustness to Distribution Shift)**.** *The likelihood ratio test for $\mathbf{H_0} : z^* \notin D_\tau$ versus $\mathbf{H_1^\tau} : z^* \in D_\tau$ depends only on $(\mu_\tau, \sigma_\tau^2)$. The distributions of batches D_t for $t \neq \tau$ do not affect the test.*

Proof. **Likelihood ratio decomposition.** The cumulative empirical mean satisfies $\hat{\mu}_t = \frac{t-1}{t}\hat{\mu}_{t-1} + \frac{1}{t}\bar{X}_t$. The likelihood ratio decomposes as:

$$\mathrm{LR} = \prod_{t=1}^{T} \frac{p(\hat{\mu}_t \mid \hat{\mu}_{t-1}, \mathbf{H_1^\tau})}{p(\hat{\mu}_t \mid \hat{\mu}_{t-1}, \mathbf{H_0})}.$$

Terms for $t \neq \tau$ cancel. For $t < \tau$, the target has not been inserted. Under both hypotheses, batch D_t is drawn from $\mathcal{N}(\mu_t, \sigma_t^2)$. The conditional distributions of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ are identical under both hypotheses, so the corresponding factor equals 1.

For $t > \tau$, batch D_t contains no target under either hypothesis. Although the effect of z^* is present in $\hat{\mu}_{t-1}$ under $\mathbf{H_1^\tau}$, conditioning on $\hat{\mu}_{t-1}$ absorbs this effect. The batch $\bar{X}_t$ is drawn from $\mathcal{N}(\mu_t, \sigma_t^2/n)$

under both hypotheses, so the conditional distributions of $\hat{\mu}_t$ given $\hat{\mu}_{t-1}$ are identical. These factors also equal 1.

The transition at $t = \tau$. The only factor that differs is at $t = \tau$, where the conditional distributions are:

$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_0} \sim \mathcal{N}\left(\frac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \frac{\mu_\tau}{\tau}, \frac{\sigma_\tau^2}{\tau^2 n}\right),$$
$$\hat{\mu}_\tau \mid \hat{\mu}_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}\left(\frac{\tau-1}{\tau}\hat{\mu}_{\tau-1} + \frac{(n-1)\mu_\tau + z^*}{\tau n}, \frac{(n-1)\sigma_\tau^2}{\tau^2 n^2}\right).$$

These expressions involve only $(\mu_\tau, \sigma_\tau^2)$. The parameters (μ_t, σ_t^2) for $t \neq \tau$ do not appear in the likelihood ratio. □

G SeMI^{Unif}

We derive the optimal test for $\mathbf{H_0}$ against $\widetilde{\mathbf{H_1}}$ where $\tau \sim \text{Uniform}(\{1, \ldots, T\})$. As before, T batches of size n are drawn from $\mathcal{D} = \mathcal{N}(\mu, \sigma^2)$ with known parameters.

Notation. We write LR_t for the likelihood ratio at time t (derived in Theorem D.1) and $\log \text{LR}_t$ for its logarithm. The batch means $\bar{X}_1, \ldots, \bar{X}_T$ are independent under both hypotheses.

G.1 Derivation of the Likelihood Ratio

Theorem G.1 (SeMI^{Unif} Likelihood Ratio)**.** *Under $\widetilde{\mathbf{H_1}}$ with $\tau \sim \text{Uniform}(\{1, \ldots, T\})$, the likelihood ratio between $\widetilde{\mathbf{H_1}}$ and $\mathbf{H_0}$ is*

$$\text{LR}_{\text{Unif}} = \frac{1}{T}\sum_{t=1}^{T} \text{LR}_t.$$

Proof. Under $\widetilde{\mathbf{H_1}}$, the insertion time satisfies $\mathbb{P}(\tau = t \mid \widetilde{\mathbf{H_1}}) = 1/T$. The marginal likelihood of the observations is:

$$p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \widetilde{\mathbf{H_1}}) = \sum_{t=1}^{T} \mathbb{P}(\tau = t) \cdot p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \tau = t, \widetilde{\mathbf{H_1}}).$$

Conditioned on $\tau = t$, the data distribution matches $\mathbf{H_1^t}$:

$$p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \tau = t, \widetilde{\mathbf{H_1}}) = p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_1^t}).$$

Dividing by the null likelihood:

$$\text{LR}_{\text{Unif}} = \frac{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \widetilde{\mathbf{H_1}})}{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_0})} = \frac{1}{T}\sum_{t=1}^{T} \frac{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_1^t})}{p(\hat{\mu}_1, \ldots, \hat{\mu}_T \mid \mathbf{H_0})} = \frac{1}{T}\sum_{t=1}^{T} \text{LR}_t.$$

□

Each LR_t depends on the data only through $\bar{X}_t$ (by the isolation property), and batch means are independent. The test statistic aggregates evidence across all time steps.

G.2 Log-Sum-Exp Form

Corollary G.2 (Log Likelihood Ratio)**.** *The log likelihood ratio is*

$$\log \text{LR}_{\text{Unif}} = \text{logsumexp}(\log \text{LR}_1, \ldots, \log \text{LR}_T) - \log T$$

where $\text{logsumexp}(x_1, \ldots, x_T) = \log\left(\sum_{t=1}^{T} e^{x_t}\right)$.

Proof. Taking the logarithm of Theorem G.1:

$$\log \mathrm{LR}_{\mathrm{Unif}} = \log\left(\frac{1}{T}\sum_{t=1}^{T}\mathrm{LR}_t\right) = \log\left(\sum_{t=1}^{T} e^{\log \mathrm{LR}_t}\right) - \log T.$$

□

The log-sum-exp function is numerically stable: $\mathrm{logsumexp}(x_1, \ldots, x_T) = m + \log\left(\sum_{t=1}^{T} e^{x_t - m}\right)$ where $m = \max_t x_t$.

G.3 Relationship to the GLR Test

Proposition G.3 (Bounds)**.** *The log likelihood ratio satisfies*

$$\max_t \log \mathrm{LR}_t - \log T \leqslant \log \mathrm{LR}_{\mathrm{Unif}} \leqslant \max_t \log \mathrm{LR}_t.$$

Proof. **Upper bound.** Since $\mathrm{LR}_t \geqslant 0$:

$$\frac{1}{T}\sum_{t=1}^{T}\mathrm{LR}_t \leqslant \max_t \mathrm{LR}_t.$$

Lower bound. The average is at least the maximum divided by T:

$$\frac{1}{T}\sum_{t=1}^{T}\mathrm{LR}_t \geqslant \frac{1}{T}\max_t \mathrm{LR}_t.$$

□

When one LR_{t^*} dominates, $\log \mathrm{LR}_{\mathrm{Unif}} \approx \log \mathrm{LR}_{t^*} - \log T$. When multiple batches contribute comparable evidence, the test accumulates their contributions.

H $\mathrm{SeMI}^{\max}$

Consider the case where T batches of size n are drawn from $\mathcal{D} = \mathcal{N}(\mu, \sigma^2)$ with known parameters (μ, σ^2), but the insertion time is unknown. The target $z^* \in \mathbb{R}$ is fixed. We test:

$\mathbf{H_0}$: All datapoints sampled i.i.d. from $\mathcal{D}$

$\overline{\mathbf{H_1}}$: One sample in D_τ replaced by z^* for some unknown $\tau \in \{1, \ldots, T\}$

Notation. We write $\bar{X}_t$ for the batch mean, $\hat{\mu}_t$ for the cumulative empirical mean, and $\log \mathrm{LR}_\tau$ for the log likelihood ratio at time τ. The test statistic for $\mathrm{SeMI}^{\max}$ is $\max_\tau \log \mathrm{LR}_\tau$. Single-batch error rates are $\alpha_0(\gamma) = P_0(\log \mathrm{LR}_1 \geqslant \gamma)$ and $\beta_0(\gamma) = P_1(\log \mathrm{LR}_1 < \gamma)$.

H.1 Independence of Log Likelihood Ratios under $\mathbf{H_0}$

Theorem H.1 (Independence under $\mathbf{H_0}$)**.** *Under* $\mathbf{H_0}$*, the statistics* $\log \mathrm{LR}_1, \ldots, \log \mathrm{LR}_T$ *are independent and identically distributed.*

Proof. From Theorem D.1, each $\log \mathrm{LR}_\tau$ depends on the data only through $\bar{X}_\tau - \mu$. The coefficients c_0, c_1, and c_2 in the quadratic form depend on n, σ^2, μ, and z^*, but not on τ.

Under $\mathbf{H_0}$, each batch D_τ is sampled independently from $\mathcal{N}(\mu, \sigma^2)^n$. The batch means $\bar{X}_1, \ldots, \bar{X}_T$ are therefore independent random variables. Moreover, each $\bar{X}_\tau \sim \mathcal{N}(\mu, \sigma^2/n)$ for all $\tau \in \{1, \ldots, T\}$. Since the log likelihood ratio is a function of $\bar{X}_\tau - \mu$ with coefficients that do not depend on τ, the statistics $\log \mathrm{LR}_1, \ldots, \log \mathrm{LR}_T$ have the same distribution.

The statistics $\log \mathrm{LR}_\tau$ are functions of independent, identically distributed random variables (the batch means), and these functions are identical across τ. Therefore, $\log \mathrm{LR}_1, \ldots, \log \mathrm{LR}_T$ are i.i.d. under $\mathbf{H_0}$. □

H.2 Type I error

Corollary H.2 ($\text{SeMI}^{\max}$ Type I Error)**.** *The Type I error of* $\text{SeMI}^{\max}$ *with threshold* γ *is*

$$\alpha_{\max}(\gamma) = 1 - [1 - \alpha_0(\gamma)]^T$$

where $\alpha_0(\gamma)$ *is the Type I error of the single-batch test from Lemma 3.3.*

Proof. The $\text{SeMI}^{\max}$ test rejects when $\max_{\tau=1,\ldots,T} \log \text{LR}_\tau \geqslant \gamma$. The Type I error is:

$$\begin{aligned}\alpha_{\max}(\gamma) &= P_0(\text{SeMI}^{\max} \geqslant \gamma) \\ &= P_0\left(\max_{\tau=1,\ldots,T} \log \text{LR}_\tau \geqslant \gamma\right) \\ &= 1 - P_0(\log \text{LR}_\tau < \gamma \text{ for all } \tau).\end{aligned}$$

By Theorem H.1, the statistics are independent under $\mathbf{H_0}$:

$$P_0(\log \text{LR}_\tau < \gamma \text{ for all } \tau) = \prod_{\tau=1}^{T} P_0(\log \text{LR}_\tau < \gamma).$$

Since all $\log \text{LR}_\tau$ have the same distribution under $\mathbf{H_0}$:

$$\prod_{\tau=1}^{T} P_0(\log \text{LR}_\tau < \gamma) = [P_0(\log \text{LR}_1 < \gamma)]^T = [1 - \alpha_0(\gamma)]^T.$$

Therefore:

$$\alpha_{\max}(\gamma) = 1 - [1 - \alpha_0(\gamma)]^T.$$

□

To achieve overall Type I error α, set the threshold γ^* such that $\alpha_0(\gamma^*) = 1 - (1 - \alpha)^{1/T}$. For small α and moderate T, this approximates the Bonferroni correction $\alpha_0 \approx \alpha/T$.

H.3 Type II error

Corollary H.3 ($\text{SeMI}^{\max}$ Type II Error)**.** *The Type II error of* $\text{SeMI}^{\max}$*, given true insertion time* τ^**, is*

$$\beta_{\max}(\gamma \mid \tau^*) = \beta_0(\gamma) \cdot [1 - \alpha_0(\gamma)]^{T-1}$$

where $\beta_0(\gamma)$ *is the Type II error from Lemma 3.3. This does not depend on* τ^**.*

Proof. Under $\mathbf{H_1^{\tau^*}}$, the target z^* is in batch τ^*. The Type II error is the probability of failing to reject $\mathbf{H_0}$.

The test fails to reject when all log likelihood ratios are below the threshold:

$$\beta_{\max}(\gamma \mid \tau^*) = P_1(\text{SeMI}^{\max} < \gamma \mid \mathbf{H_1^{\tau^*}}) = P_1\left(\max_\tau \log \text{LR}_\tau < \gamma\right).$$

Separating the true insertion time from others:

$$P_1\left(\max_\tau \log \text{LR}_\tau < \gamma\right) = P_1\left(\log \text{LR}_{\tau^*} < \gamma \text{ and } \max_{\tau \neq \tau^*} \log \text{LR}_\tau < \gamma\right).$$

Under $\mathbf{H_1^{\tau^*}}$, the statistic $\log \text{LR}_{\tau^*}$ depends only on batch D_{τ^*}, while $\log \text{LR}_\tau$ for $\tau \neq \tau^*$ depends only on batch D_τ. Since the batches are independent, these statistics are independent:

$$P_1\left(\log \text{LR}_{\tau^*} < \gamma \text{ and } \max_{\tau \neq \tau^*} \log \text{LR}_\tau < \gamma\right) = P_1(\log \text{LR}_{\tau^*} < \gamma) \cdot P_1\left(\max_{\tau \neq \tau^*} \log \text{LR}_\tau < \gamma\right).$$

The first factor is $\beta_0(\gamma)$, the Type II error of the single-batch test. For batches $\tau \neq \tau^*$, no target is present, so $\log \text{LR}_\tau$ has the null distribution. By the same argument as in Corollary H.2:

$$P_1\left(\max_{\tau \neq \tau^*} \log \text{LR}_\tau < \gamma\right) = [1 - \alpha_0(\gamma)]^{T-1}.$$

Therefore:

$$\beta_{\max}(\gamma \mid \tau^*) = \beta_0(\gamma) \cdot [1 - \alpha_0(\gamma)]^{T-1}.$$

The result does not depend on τ^* because both $\alpha_0(\gamma)$ and $\beta_0(\gamma)$ are independent of the insertion time. □

The Type II error decomposes as a product: $\beta_0(\gamma)$ is the probability that the test fails to detect the target in the correct batch, and $[1 - \alpha_0(\gamma)]^{T-1}$ is the probability that none of the other batches trigger a false detection. Independence from τ^* follows from the identical distribution of batch means.

I SGD Analysis

We consider the SGD setting where we observe model parameters $\theta_0, \theta_1, \ldots, \theta_T$ with $\theta_t \in \mathbb{R}^d$. At each step t, the update is $\theta_t = \theta_{t-1} - \eta_t g_t$ where $\eta_t > 0$ is the learning rate and g_t is the batch gradient on dataset D_t of size n_t. The target z^* lies in the data domain $\mathcal{X}$ and the insertion time τ is known. We test:

$$\mathbf{H_0} : \text{All samples in } D_1, \ldots, D_T \text{ are drawn i.i.d. from } \mathcal{D}$$
$$\mathbf{H_1^\tau} : \text{One sample in } D_\tau \text{ is replaced by } z^*$$

We assume that, conditioned on θ_{t-1}, the batch gradient g_t is Gaussian with mean $\mu_g(\theta_{t-1})$ and covariance $\Sigma_g(\theta_{t-1})/n_t$. This approximation is justified by the central limit theorem when the batch aggregates gradients from many data points.

Notation. The loss function on data point x with parameter θ is denoted by $\ell(\theta; x)$. The batch gradient is $g_t = \frac{1}{n_t} \sum_{x \in D_t} \nabla_\theta \ell(\theta_{t-1}; x)$. We denote by $\mu_g(\theta)$ and $\Sigma_g(\theta)$ the mean and covariance of the batch gradient when scaled by $\sqrt{n_t}$, i.e., $g_t \mid \theta_{t-1} \sim \mathcal{N}(\mu_g(\theta_{t-1}), \Sigma_g(\theta_{t-1})/n_t)$.

I.1 Conditional Distributions

Lemma I.1 (Parameter Distribution Under $\mathbf{H_0}$)**.** *Under $\mathbf{H_0}$, the conditional distribution of θ_t given θ_{t-1} is:*

$$\theta_t \mid \theta_{t-1}, \mathbf{H_0} \sim \mathcal{N}\left(\theta_{t-1} - \eta_t \mu_g(\theta_{t-1}), \frac{\eta_t^2 \Sigma_g(\theta_{t-1})}{n_t}\right).$$

Lemma I.2 (Parameter Distribution Under $\mathbf{H_1^\tau}$)**.** *Under $\mathbf{H_1^\tau}$, the conditional distribution of θ_τ given $\theta_{\tau-1}$ is:*

$$\theta_\tau \mid \theta_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}\left(\theta_{\tau-1} - \eta_\tau \cdot \frac{(n_\tau - 1)\mu_g(\theta_{\tau-1}) + \nabla_\theta \ell(\theta_{\tau-1}; z^*)}{n_\tau}, \frac{\eta_\tau^2 (n_\tau - 1)\Sigma_g(\theta_{\tau-1})}{n_\tau^2}\right).$$

For $t \neq \tau$, the conditional distribution is identical to that under $\mathbf{H_0}$.

Proof. **Batch gradient under $\mathbf{H_1^\tau}$.** Batch D_τ contains $(n_\tau - 1)$ samples from $\mathcal{D}$ and the fixed target z^*. The batch gradient is:

$$g_\tau = \frac{1}{n_\tau}\left[\sum_{i=1}^{n_\tau - 1} \nabla_\theta \ell(\theta_{\tau-1}; x_i) + \nabla_\theta \ell(\theta_{\tau-1}; z^*)\right].$$

Under the Gaussian gradient assumption, the sum of $(n_\tau - 1)$ i.i.d. Gaussian gradients has mean $(n_\tau - 1)\mu_g(\theta_{\tau-1})$ and covariance $(n_\tau - 1)\Sigma_g(\theta_{\tau-1})$. Adding the fixed target gradient shifts the mean.

The result follows from $\theta_\tau = \theta_{\tau-1} - \eta_\tau g_\tau$. □

I.2 Likelihood Ratio

Theorem I.3 (SGD Likelihood Ratio Simplification)**.** *The likelihood ratio for testing $\mathbf{H_0}$ against $\mathbf{H_1^\tau}$ based on $(\theta_0, \ldots, \theta_T)$ satisfies:*

$$\frac{p(\theta_0, \ldots, \theta_T \mid \mathbf{H_1^\tau})}{p(\theta_0, \ldots, \theta_T \mid \mathbf{H_0})} = \frac{p(\theta_\tau \mid \theta_{\tau-1}, \mathbf{H_1^\tau})}{p(\theta_\tau \mid \theta_{\tau-1}, \mathbf{H_0})}.$$

Proof. By the chain rule,

$$\frac{p(\theta_0, \ldots, \theta_T \mid H)}{p(\theta_0)} = \prod_{t=1}^{T} p(\theta_t \mid \theta_0, \ldots, \theta_{t-1}, H).$$

The update $\theta_t = \theta_{t-1} - \eta_t g_t$ implies that θ_t depends on $(\theta_0, \ldots, \theta_{t-1})$ only through θ_{t-1}, since g_t depends on θ_{t-1} and the samples in D_t. Thus

$$p(\theta_t \mid \theta_0, \ldots, \theta_{t-1}, H) = p(\theta_t \mid \theta_{t-1}, H).$$

The likelihood ratio becomes

$$\mathrm{LR} = \prod_{t=1}^{T} \frac{p(\theta_t \mid \theta_{t-1}, \mathbf{H_1^\tau})}{p(\theta_t \mid \theta_{t-1}, \mathbf{H_0})}.$$

For $t < \tau$, the target has not been inserted, so both conditional distributions are identical and the ratio equals 1. For $t > \tau$, batch D_t contains no target under either hypothesis, so conditioned on θ_{t-1} the distributions coincide and the ratio equals 1. Only the term at $t = \tau$ contributes. □

Theorem I.4 (SGD Log Likelihood Ratio)**.** *Define* $\delta_g = \nabla_\theta \ell(\theta_{\tau-1}; z^*) - \mu_g(\theta_{\tau-1})$ *and* $N = \theta_\tau - \theta_{\tau-1} + \eta_\tau \mu_g(\theta_{\tau-1})$*. The log likelihood ratio is:*

$$\log \mathrm{LR} = -\frac{d}{2} \log\left(\frac{n_\tau - 1}{n_\tau}\right) - \frac{n_\tau}{2(n_\tau - 1)\eta_\tau^2} N^\top \Sigma_g^{-1} N - \frac{n_\tau}{(n_\tau - 1)\eta_\tau} N^\top \Sigma_g^{-1} \delta_g - \frac{m^*}{2(n_\tau - 1)}$$

where $m^* = \delta_g^\top \Sigma_g^{-1} \delta_g$ *is the Mahalanobis distance of the target gradient from the population mean. We denote this sequential test for SGD as* $\mathrm{SeMI}^{\mathrm{SGD}}$.

Proof. **Conditional parameters.** Let $m_0 = \theta_{\tau-1} - \eta_\tau \mu_g(\theta_{\tau-1})$ and $m_1 = \theta_{\tau-1} - \eta_\tau \frac{(n_\tau - 1)\mu_g(\theta_{\tau-1}) + \nabla_\theta \ell(\theta_{\tau-1}; z^*)}{n_\tau}$. The covariances are $V_0 = \frac{\eta_\tau^2 \Sigma_g}{n_\tau}$ and $V_1 = \frac{\eta_\tau^2 (n_\tau - 1)\Sigma_g}{n_\tau^2}$.

Mean difference.

$$m_1 - m_0 = -\frac{\eta_\tau}{n_\tau}(\nabla_\theta \ell(\theta_{\tau-1}; z^*) - \mu_g(\theta_{\tau-1})) = -\frac{\eta_\tau \delta_g}{n_\tau}.$$

Log likelihood ratio. With $N = \theta_\tau - m_0$, expand the quadratic terms and substitute $V_0^{-1} = \frac{n_\tau}{\eta_\tau^2} \Sigma_g^{-1}$. Collecting terms yields the stated result. □

Theorem I.5 (Distribution of the Log Likelihood Ratio)**.** *Under* $\mathbf{H_0}$*:*

$$N \mid \theta_{\tau-1}, \mathbf{H_0} \sim \mathcal{N}\left(0, \frac{\eta_\tau^2 \Sigma_g}{n_\tau}\right).$$

Under $\mathbf{H_1^\tau}$*:*

$$N \mid \theta_{\tau-1}, \mathbf{H_1^\tau} \sim \mathcal{N}\left(-\frac{\eta_\tau \delta_g}{n_\tau}, \frac{\eta_\tau^2 (n_\tau - 1)\Sigma_g}{n_\tau^2}\right).$$

Proof. Under $\mathbf{H_0}$, $\theta_\tau \mid \theta_{\tau-1} \sim \mathcal{N}(m_0, V_0)$ with $m_0 = \theta_{\tau-1} - \eta_\tau \mu_g(\theta_{\tau-1})$. Since $N = \theta_\tau - m_0$, we have $N \sim \mathcal{N}(0, V_0)$.

Under $\mathbf{H_1^\tau}$, $\theta_\tau \mid \theta_{\tau-1} \sim \mathcal{N}(m_1, V_1)$. Thus $N = \theta_\tau - m_0 \sim \mathcal{N}(m_1 - m_0, V_1)$. Recalling that $m_1 - m_0 = -\frac{\eta_\tau \delta_g}{n_\tau}$, we have the stated mean and covariance. □

I.3 Application to Linear Regression

For linear regression with squared loss $\ell(\theta;(x,y)) = \frac{1}{2}(y-\theta^\top x)^2$ and Gaussian covariates $x \sim \mathcal{N}(0,\Sigma_x)$:

Lemma I.6 (Mean Gradient). *Let $\Delta = \theta - \theta^*$. The mean gradient is $\mu_g(\theta) = \Sigma_x \Delta$.*

Proof. Let $\nabla_\theta \ell(\theta;(x,y)) = (\theta^\top x - y)x = (\theta^\top x - \theta^{*\top} x - \varepsilon)x = (\Delta^\top x - \varepsilon)x$. Taking expectations:

$$\mu_g(\theta) = \mathbb{E}[(\Delta^\top x)x] - \mathbb{E}[\varepsilon]\mathbb{E}[x] = \mathbb{E}[xx^\top]\Delta = \Sigma_x \Delta.$$

□

Lemma I.7 (Gradient Covariance). *The covariance of the per-sample gradient is:*

$$\Sigma_g(\theta) = (\sigma_\varepsilon^2 + \Delta^\top \Sigma_x \Delta)\Sigma_x + \Sigma_x \Delta\Delta^\top \Sigma_x.$$

Proof. The second moment is $\mathbb{E}[gg^\top] = \mathbb{E}[(\Delta^\top x - \varepsilon)^2 xx^\top]$. Since $\varepsilon \perp x$ and $\mathbb{E}[\varepsilon] = 0$:

$$\mathbb{E}[gg^\top] = \mathbb{E}[(\Delta^\top x)^2 xx^\top] + \sigma_\varepsilon^2 \Sigma_x.$$

Using Isserlis' theorem for the fourth moment of Gaussian x:

$$\mathbb{E}[(\Delta^\top x)^2 xx^\top] = (\Delta^\top \Sigma_x \Delta)\Sigma_x + 2\Sigma_x \Delta\Delta^\top \Sigma_x.$$

The covariance is $\Sigma_g = \mathbb{E}[gg^\top] - \mu_g \mu_g^\top$. Subtracting $\mu_g \mu_g^\top = \Sigma_x \Delta\Delta^\top \Sigma_x$ yields the result. □

Corollary I.8 (Gradient Statistics at Optimum). *At $\theta = \theta^*$, we have $\Delta = 0$, so $\mu_g(\theta^*) = 0$ and $\Sigma_g(\theta^*) = \sigma_\varepsilon^2 \Sigma_x$.*

Lemma I.9 (Target Gradient Deviation). *At $\theta = \theta^*$, the target gradient is $\nabla_\theta \ell(\theta^*;(x^*,y^*)) = -\varepsilon^* x^*$ where ε^* is the target residual. The deviation from the population mean is $\delta_g = -\varepsilon^* x^*$.*

Proof. At optimum, $\nabla_\theta \ell = (\theta^{*\top} x^* - y^*)x^* = -\varepsilon^* x^*$. Since $\mu_g(\theta^*) = 0$, $\delta_g = -\varepsilon^* x^*$. □

Theorem I.10 (Detectability at Optimum). *At $\theta = \theta^*$, the Mahalanobis distance of the target gradient is:*

$$m^* = \frac{(\varepsilon^*)^2}{\sigma_\varepsilon^2} \cdot (x^*)^\top \Sigma_x^{-1} x^*.$$

This factors as product of a label outlier score $(\varepsilon^*/\sigma_\varepsilon)^2$ *and a* feature leverage score $(x^*)^\top \Sigma_x^{-1} x^*$.

Proof. Substituting into $m^* = \delta_g^\top \Sigma_g(\theta^*)^{-1} \delta_g$:

$$m^* = (-\varepsilon^* x^*)^\top (\sigma_\varepsilon^2 \Sigma_x)^{-1} (-\varepsilon^* x^*) = \frac{(\varepsilon^*)^2}{\sigma_\varepsilon^2} (x^*)^\top \Sigma_x^{-1} x^*.$$

□

Samples with large residuals (relative to population noise) and unusual feature vectors (relative to population covariance) are most detectable.

J Additional Experiments

This appendix provides additional experimental results supporting the main paper. Section J.1 reports the empirical-mean illustrations under the disjoint-batch protocol that exactly matches the SeMI* assumptions. Section J.2 extends the unknown-τ comparison. Section J.3 validates SeMI^{SGD} on linear regression. Section J.4 reports the multivariate Gaussian setting. Section J.5 reports a complementary DP-SGD pretrain-then-finetune protocol on three datasets. Section J.6 quantifies the gap between averaging over τ and selecting the best τ post hoc.

Compute. The main DP-SGD audit on Fashion-MNIST and the three pretrain-then-finetune audits ran on a single NVIDIA GeForce RTX 2080 Ti: approximately 48 hours for the main audit, and approximately 5 hours each for the Fashion-MNIST, CIFAR-10, and Purchase-100 fine-tuning runs. The empirical-mean illustrations (Appendices J.1–J.2), the SGD validation on linear regression (Appendix J.3), and the multivariate experiment (App. J.4) run on CPU in minutes.

J.1 Disjoint-Batch Protocol

The numerical illustrations in Section 3.1 use a subsampled-with-replacement protocol that departs from the i.i.d. disjoint-batch assumption underlying SeMI^*. Here we report the matching figures under the disjoint protocol that exactly satisfies the test assumptions: $T = 10$ independent batches of size $n = 10$ drawn i.i.d. from $\mathcal{N}(0, 1)$, target z^* inserted at $\tau = 5$ with Mahalanobis distance 5, $R = 50\,000$ runs.

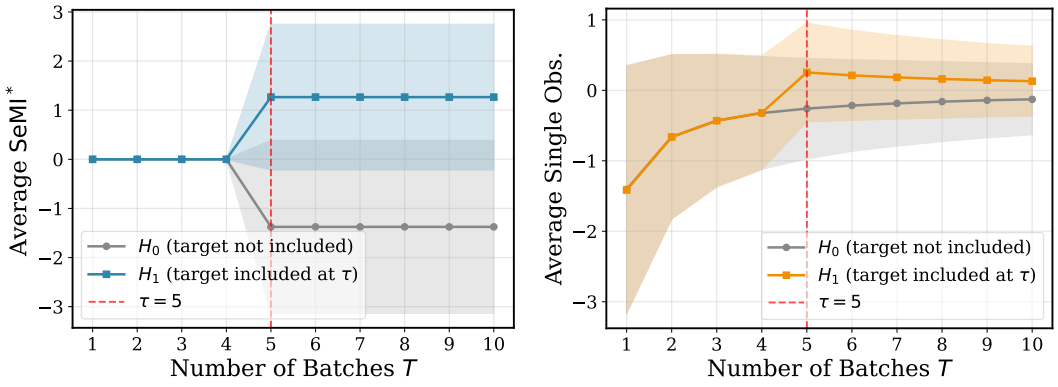

(a) SeMI^* statistic (b) Final Obs. statistic

Figure 7: Log-likelihood ratio under the disjoint-batch protocol ($n = 10$, $\tau = 5$).

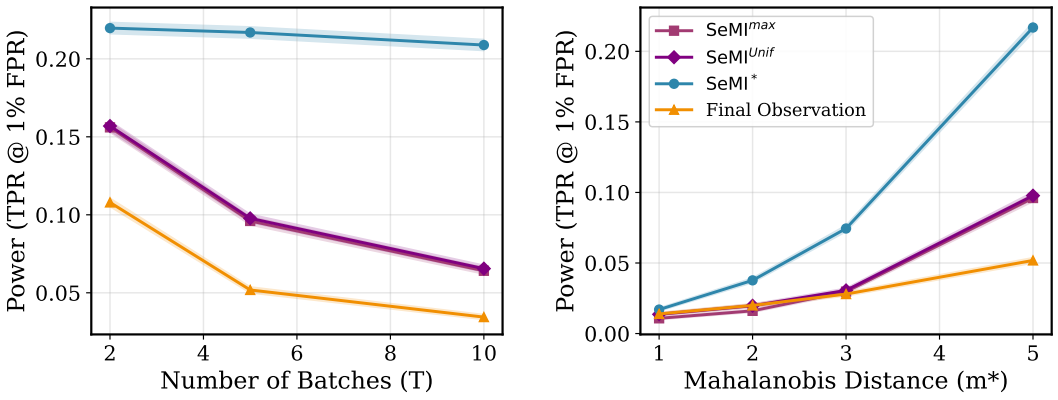

(a) Power vs T (b) Power vs z^* strength

Figure 8: Power comparison under the disjoint-batch protocol ($n = 10$, TPR at 1% FPR). The qualitative behavior matches the subsampled-with-replacement protocol used in the main text.

J.2 Unknown Insertion Time

This section provides additional results for the sequential tests when the insertion time τ is unknown. We use $n = 10$, $\mu = 0$, $\sigma = 1$, and target z^* with Mahalanobis distance 5. Power comparisons as a function of T and signal strength appear in the main text (Figure 3).

Figure 9 shows ROC curves for different values of T. As T increases, the gap between SeMI^* and the other methods widens, while $\mathrm{SeMI}^{\mathrm{Unif}}$ and $\mathrm{SeMI}^{\mathrm{max}}$ maintain an advantage over Final Observation.

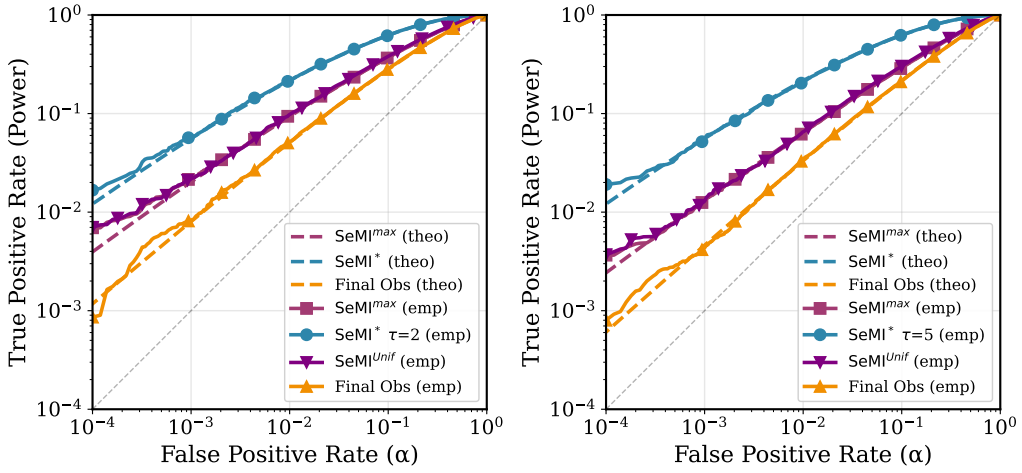

(a) $T = 5$ (b) $T = 10$

Figure 9: ROC curves comparing SeMI^*, $\text{SeMI}^{\text{Unif}}$, $\text{SeMI}^{\max}$, and Final Observation for different T.

J.3 SGD Validation

We validate the extension to SGD (Section 4) using linear regression with $d = 5$ features, $T = 10$ updates, batch size $n = 50$, and learning rate $\eta = 0.05$.

Figure 10(a) shows ROC curves for insertion at $\tau = 5$. SeMI^{SGD} outperforms the Delta and Back-Front heuristics from Jagielski et al. (2023).

Figure 10(b) shows power as a function of insertion time τ. The Back-Front statistic depends only on the endpoints θ_0 and θ_T, so its power is approximately invariant to where in the trajectory the target is inserted. The Delta statistic shows decreasing power as τ increases: early in training, gradient magnitudes are larger, producing stronger signals. SeMI^{SGD} maintains high power across all τ.

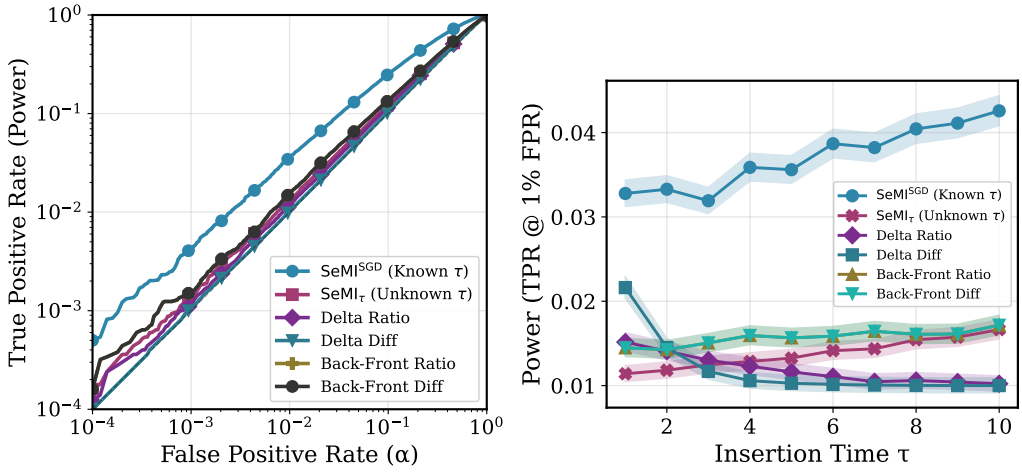

(a) ROC Comparison ($\tau = 5$) (b) Power vs Insertion Time τ

Figure 10: SGD validation on linear regression. (a) ROC curves. (b) Power vs τ. SeMI^{SGD} maintains power across τ in this linear-regression setting; on the deep-model audit of Section 5, audit tightness instead grows with τ (Observation 4).

J.4 Multivariate Setting

We test the extension to d-dimensional Gaussian data (Appendix E). Parameters: $n = 10$, $T = 5$, $\tau = 3$, target z^* with Mahalanobis distance 3 from the mean.

Figure 11(a) displays the decision boundary in $d = 2$. Figure 11(b) shows that power remains constant as dimension d increases from 2 to 50. For a fixed Mahalanobis distance, the test statistic depends on the projection onto the target direction, making power independent of ambient dimension.

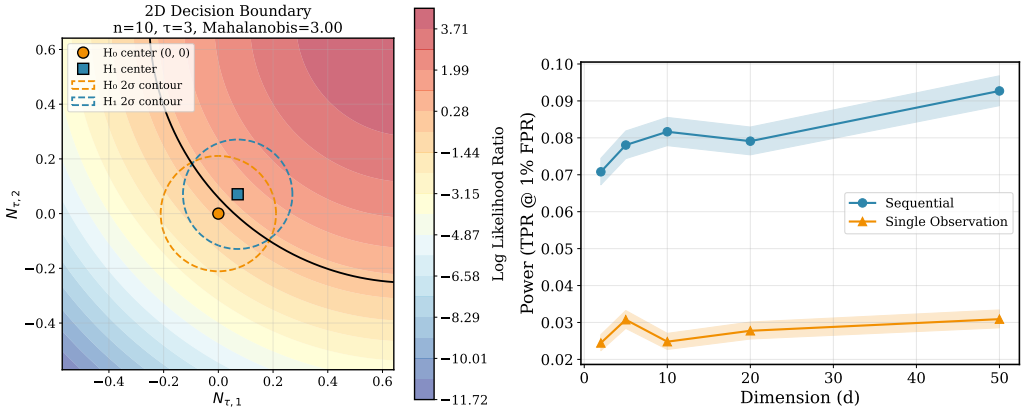

(a) Decision Boundary ($d = 2$) (b) Power vs Dimension

Figure 11: Multivariate setting. (a) Decision boundary in 2D. (b) Power vs dimension d.

J.5 DP-SGD Fine-tuning Instantiation

Setup. This is a separate protocol from the disjoint-phase audit of Section 5: instead of splitting one DP-SGD run into T phases on disjoint subsets, we pretrain on a public split with SGD and then run T DP-SGD updates on a private split, averaging audit bounds over $\tau \in \{1, \ldots, T\}$. We use batch size 64, $T = 10$ updates, and $\delta = 10^{-4}$. The auditor observes the parameter trajectory $(\theta_0, \ldots, \theta_T)$ and knows the insertion time τ. For each audit, the target z^* is selected from a candidate pool by taking the sample with the highest gradient Mahalanobis distance, estimated from a held-out reference set. We average results over $\tau \in \{1, \ldots, T\}$; Appendix J.6 shows that selecting the best τ post hoc yields tighter bounds, and how to choose τ a priori for a given architecture remains open. Three architectures are tested: multinomial logistic regression on Fashion-MNIST (Xiao et al., 2017), a frozen pretrained VGG-16 (Simonyan and Zisserman, 2014) with a single trainable layer on CIFAR-10 (Krizhevsky et al., 2009), and a fully connected network with 128 hidden units on Purchase-100 (Shokri et al., 2017).

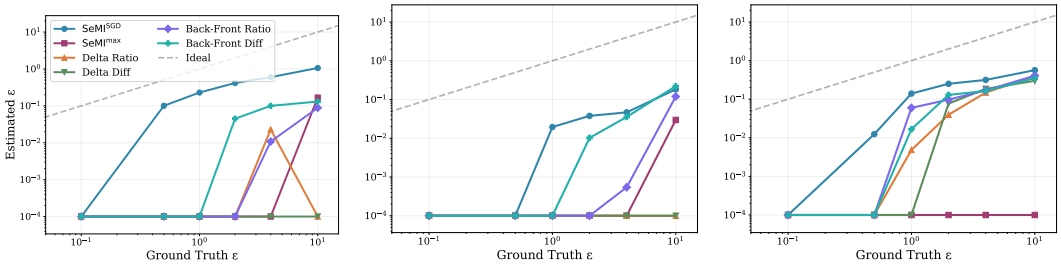

(a) Fashion-MNIST (b) CIFAR-10 (c) Purchase-100

Figure 12: Private fine-tuning with $T = 10$ DP-SGD updates and known τ. Estimated ε lower bound versus ground-truth ε, averaged over $\tau \in \{1, \ldots, T\}$. Closer to the diagonal is tighter. SeMI^{SGD} yields tighter lower bounds than loss-based heuristics across the privacy range.

ε lower bounds. Figure 12 shows the audit lower bound $\underline{\varepsilon}$ versus ground-truth ε on the three datasets. On Fashion-MNIST, SeMI^{SGD} produces tighter lower bounds than every baseline across the privacy

range. On CIFAR-10 and Purchase-100, SeMI^{SGD} matches or beats the baselines, with the clearest gains in low-privacy regimes.

ROC curves. Figure 13 reports ROC curves on Fashion-MNIST at two privacy levels. Without DP noise, SeMI^{SGD} and SeMI^{max} both achieve near-perfect discrimination (AUC ≈ 1.0), while the heuristic baselines perform substantially worse (Back-Front AUC ≈ 0.80, Delta AUC ≈ 0.65). At $\varepsilon = 1.0$, all methods degrade due to the DP noise, but SeMI^{SGD} maintains an advantage (AUC $= 0.60$) over baselines (AUC ≈ 0.51–0.53); the gap is most pronounced at low false positive rates, which drive audit tightness.

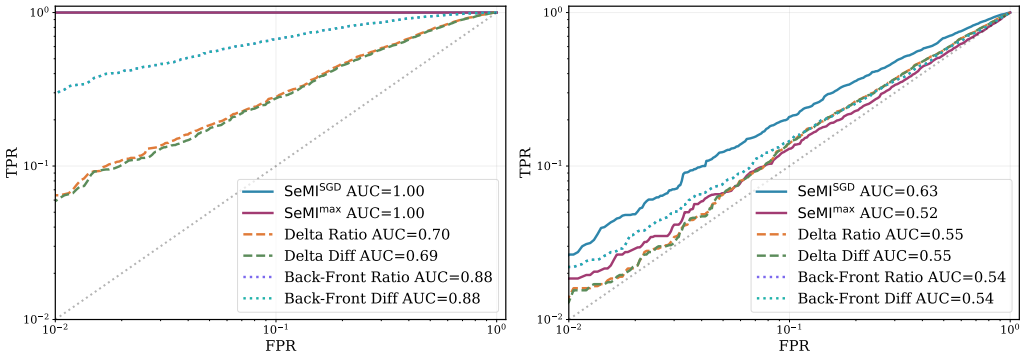

(a) Non-private ($\varepsilon = \infty$)

(b) Private ($\varepsilon = 1.0$)

Figure 13: ROC curves for Fashion-MNIST pretrain-then-finetune with DP-SGD. (a) Without DP noise, SeMI^{SGD} achieves near-perfect discrimination. (b) With DP-SGD ($\varepsilon = 1.0$), SeMI^{SGD} maintains an advantage over heuristic baselines. Cut at FPR $= 0.01$.

J.6 Effect of Insertion Time Selection

Appendix J.5 reports privacy lower bounds averaged over insertion times $\tau \in \{1, \ldots, T\}$. Here we compare this averaging strategy against selecting the best τ post hoc, i.e., reporting the tightest bound across all insertion times.

Figure 14 shows results for Fashion-MNIST. Solid lines show the averaged estimates; dashed lines show the maximum (best τ) estimates. For SeMI^{SGD}, selecting the best τ yields substantially tighter bounds, particularly at moderate privacy levels ($\varepsilon \in [0.5, 2]$). The gap between Average and Max indicates that some insertion times are more favorable for auditing than others.

This observation suggests that an auditor who can choose when to insert the target, or who audits at multiple insertion times and reports the best result, can obtain tighter privacy lower bounds.

J.7 Smaller Privacy Budgets

We rerun the Fashion-MNIST audit of Section 5 at $\varepsilon \in \{1, 2, 4\}$, keeping the canary, the protocol and the insertion time of the main text ($T = 5$, $\tau = 5$). Table 2 reports the AUC and a single $\hat{\varepsilon}$ computed over all audit rounds, in place of the subsampled violins of Figure 6. The ranking of Section 5 holds at every ε, with the sequential attacks ahead of the loss heuristics and LiRA last, since it sees only the final model. Smaller ε makes the methods harder to separate, as the DP noise drives every attack towards chance and $\hat{\varepsilon}$ towards zero.

J.8 Gaussianity of Batch Gradients

The assumption is on the batch-mean gradient, which is closer to Gaussian at the large batch sizes typical of DP training. To test it, we freeze $\theta_{\tau-1}$ from one audited run ($T = 5$, $\tau = 5$), draw 4000 batch-gradient means per dataset under SGD and DP-SGD, and test normality along the direction the score weights, $\Sigma_g^{-1}\delta_g$, and along the four leading principal directions. Normality cannot be tested directly in dimension $d \sim 10^4$ from 4000 samples, so both tests in Table 3 are on one-dimensional projections. The first column reports a Shapiro–Wilk p-value, and the second counts rejections at level

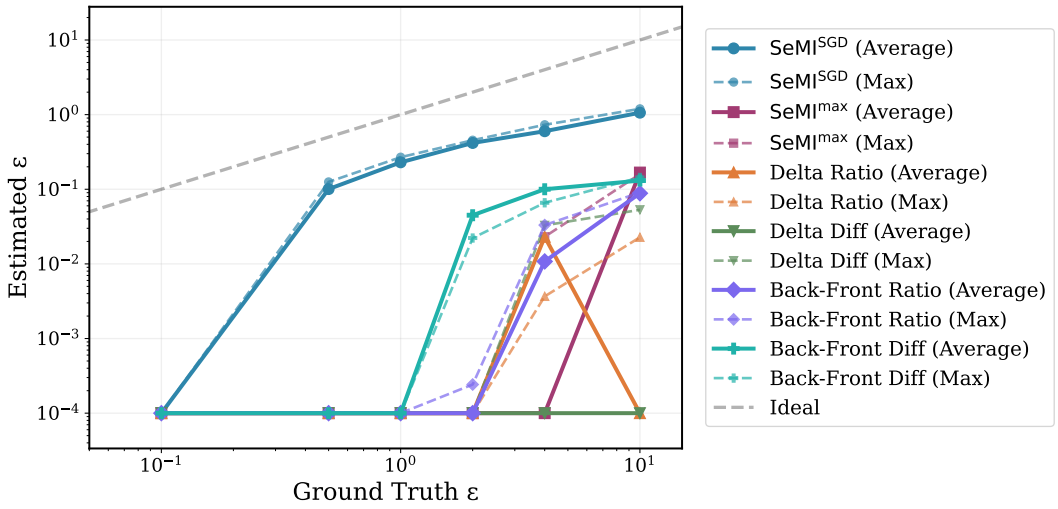

Figure 14: Effect of insertion time selection on Fashion-MNIST. Solid lines: averaged over τ. Dashed lines: best τ (post hoc selection). Selecting the optimal insertion time yields tighter privacy lower bounds.

Table 2: AUC / $\hat{\varepsilon}$ on Fashion-MNIST at $T = 5$, $\tau = 5$, $\delta = 10^{-4}$.

Attack	$\varepsilon = 1$	$\varepsilon = 2$	$\varepsilon = 4$	$\varepsilon = 8$
SeMI^{SGD}	0.535 / 0.06	0.549 / 0.11	0.595 / 0.30	0.623 / 0.49
SeMI$^{\mathrm{SGD}}_{\mathrm{max}}$	0.537 / 0.02	0.538 / 0.09	0.592 / 0.31	0.608 / 0.41
Loss Diff	0.523 / 0.00	0.527 / 0.00	0.529 / 0.02	0.534 / 0.10
Loss Ratio	0.521 / 0.00	0.526 / 0.00	0.530 / 0.00	0.533 / 0.07
LiRA	0.515 / 0.00	0.519 / 0.00	0.501 / 0.00	0.509 / 0.00

5% among the four principal directions. Gaussianity is accepted everywhere except Fashion-MNIST under plain SGD, where the projection has skew $+0.11$ and excess kurtosis -0.09. Every test accepts normality under DP-SGD, the regime we audit.

Table 3: Normality tests on batch-gradient means at $\theta_{\tau-1}$ ($T = 5$, $\tau = 5$, 4000 samples).

Data / regime	$\Sigma_g^{-1}\delta_g$ direction, p-value	PCA directions rejected
Fashion-MNIST SGD	0.008	0/4
Fashion-MNIST DP-SGD	0.35	0/4
CIFAR-10 SGD	0.40	0/4
CIFAR-10 DP-SGD	0.70	0/4
Purchase-100 SGD	0.83	0/4
Purchase-100 DP-SGD	0.91	0/4

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The three contributions stated in the abstract and Section 1 map to Sections 2, 3, 4, and 5. Assumptions and scope are stated where claims are made.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 6 discusses the Gaussian batch-gradient assumption, and Observation 4 in Section 5 the open question of choosing τ a priori. Section 5 reports regimes where gains are largest and smallest.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Each theorem and lemma states its assumptions in the body. Full proofs are cross-referenced and given in the appendix (optimal-test, uniform-prior, unknown-τ, multivariate, SGD, audit lemma).

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Section 5 reports the training pipeline, hyperparameters, and audit-protocol parameters; Appendix J extends this to the fine-tuning experiments. All datasets are public.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The code with reproduction scripts is provided as supplemental material. All datasets used are publicly available.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Optimizer, learning rate, clipping norm, batch size, number of steps, phase decomposition, splits, and audit-protocol parameters are reported in Section 5; architecture-specific details are in Appendix J.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: $\hat{\varepsilon}$ uses Clopper–Pearson upper bounds on Type-I/II errors over $R = 1900$ rounds with a union bound across the calibrated grid at overall confidence $1 - \xi = 0.98$ (Section 5, Appendix C). Violins in Figure 6 come from subsampling 90% of the rounds without replacement.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.

- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Hardware (NVIDIA GeForce RTX 2080 Ti) and approximate runtimes per experiment are reported in Appendix J.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The work uses public benchmark datasets, involves no human subjects, and contributes to privacy auditing of differentially private training, in line with the Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Tighter privacy audits improve empirical accountability of DP training (positive). The same attacks could amplify membership inference on undefended models; the focus on auditing mechanisms with formal ε guarantees bounds this risk.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper releases no pretrained models, generative models, or scraped datasets. Released artifacts are auditing scripts on public benchmarks.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Datasets (Fashion-MNIST, CIFAR-10, Purchase-100) and architectures (VGG-16) are cited at first use, as are baseline attacks (LiRA, Delta, Back-Front). Licenses are listed in the supplemental README.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes]

    Justification: The supplemental code includes a README with reproduction instructions.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing or human-subjects research.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

Justification: The paper does not involve human-subjects research.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: LLMs are not part of the core methodology.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.