NONASYMPTOTIC ANYTIME-VALID INFERENCE FOR QUANTILES UNDER LOCAL DIFFERENTIAL PRIVACY

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose a framework for sequential inference on population quantiles under local differential privacy (LDP). Specifically, we construct a nonasymptotic, anytime-valid confidence sequence by sequentially testing candidate quantiles with betting e-processes and permanently eliminating those rejected at level α. A ternary randomized response mechanism based on adaptively chosen query intervals provides the LDP guarantee. The proposed algorithm updates the confidence sequence in a single pass by taking the closed convex hull of the surviving candidates, while a nested-lattice implementation reduces memory requirements. We establish distribution-free, nonasymptotic coverage of at least $1-\alpha$ uniformly over time, ensuring valid inference at data-dependent stopping times. Under suitable regularity conditions, the interval width contracts at the iterated-logarithm rate $\mathcal{O}(\sqrt{\log\log t/t})$. Numerical experiments demonstrate the finite-sample validity and practical effectiveness of our algorithm.

1 INTRODUCTION

Quantiles are fundamental statistical summaries that describe a distribution at every level, from its center to its tails (Koenker, 2005). This makes them particularly useful for skewed and heavy-tailed data. In many applications, data are collected sequentially, and storage constraints make retaining and repeatedly accessing the full sample infeasible. In addition, collecting and analyzing such streams poses significant risks to individual privacy. For example, smartphone power consumption measurements collected over time have been shown to reveal users' locations (Michalevsky et al., 2015). These lead to a central question: how can we quantify uncertainty about a population quantile throughout a single-pass data stream while protecting individual privacy?

Differential privacy (DP) has become a standard framework for rigorous privacy protection (Dwork et al., 2006; Dwork & Roth, 2014). It requires the output distribution of a randomized mechanism to remain nearly unchanged when one individual's data are modified, making it difficult for an adversary to distinguish between the original and modified datasets from the released output. DP has two main models: central DP relies on a trusted curator who holds the raw data, whereas local DP (LDP) privatizes each observation before it leaves the user (Warner, 1965; Duchi et al., 2018). LDP is thus well suited to streams collected from personal devices and has been deployed in large-scale telemetry (Erlingsson et al., 2014; Ding et al., 2017; Tang et al., 2017). Accordingly, we focus on quantile inference under LDP.

Locally private quantile inference has attracted growing attention (Liu et al., 2023; 2024; Cai et al., 2025; 2026). In particular, Liu et al. (2023) construct self-normalized asymptotic confidence intervals (CIs) for online quantile estimation. Such CIs provide fixed-time coverage guarantees, which do not ensure validity under continuous monitoring or data-dependent stopping (Johari et al., 2022). Confidence sequences (CSs) resolve this issue by covering the parameter uniformly over time (Darling & Robbins, 1967; Howard et al., 2021; Ramdas et al., 2023). Without privacy, Chen et al. (2023) derive nonasymptotic confidence bounds for recursive stochastic gradient descent (SGD) quantile estimators under density smoothness, while Howard & Ramdas (2022) construct order-statistics-based quantile CSs through empirical-CDF inversion. Neither method extends directly to LDP: privatization changes the gradient noise underlying the SGD bounds, and the order statistics required by the latter method are not directly available from privatized reports. Under LDP, Waudby-Smith et al. (2023) establish nonasymptotic CSs for means of observations bounded in a known interval, by

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

combining nonparametric randomized response with Hoeffding-type and variance-adaptive bounds. These mean CSs do not apply directly to quantiles. A quantile is tested by comparing observations with a threshold, which may be chosen adaptively to locate the unknown target. The comparison means then vary over time, whereas their sequentially interactive CSs require a fixed mean. Most closely related, Cai et al. (2025) extend the asymptotic mean CSs of Waudby-Smith et al. (2024) to quantiles under LDP, using a Gaussian approximation to parallel SGD. In practice, inference may stop before this approximation is accurate, potentially leading to undercoverage. We therefore seek nonasymptotic simultaneous coverage for quantiles under LDP, without a known bounded data range. Table 1 summarizes the distinctions.

Table 1: Comparison of representative CS guarantees.

Work	Target	Privacy	Guarantee
Howard & Ramdas (2022)	Quantile	None	Nonasymptotic CS
Cai et al. (2025)	Quantile	ε-LDP	Asymptotic CS
Waudby-Smith et al. (2023)	Bounded mean	ε-LDP	Nonasymptotic CS
This work	Quantile	ε-LDP	**Nonasymptotic CS**

In this paper, we construct a nonasymptotic LDP quantile CS by combining sequential testing with permanent candidate elimination using only privatized reports (Figure 1). To conduct these tests, an online tracker adaptively selects a query interval, and each user provides a single ternary randomized response encoding their observation's position relative to its endpoints (Warner, 1965; Waudby-Smith et al., 2023). Each report updates betting e-processes (Waudby-Smith & Ramdas, 2024; Ramdas et al., 2023) that accumulate evidence against candidate quantiles. Candidates whose evidence reaches $1/\alpha$ are permanently excluded, and the closed convex hull of the remaining candidates forms the CS. The e-process construction controls the probability of ever excluding the true quantile, ensuring nonasymptotic simultaneous coverage. The tracker guides finite-interval acquisition without a prespecified interval containing the quantile, then refines queries as the CS contracts. The main technical challenges are to obtain a finite interval from arbitrary initialization and control peak storage while preserving exact candidate-wise evidence updates. To this end, we incorporate a prepaid capital reserve to mitigate early betting losses and use a nested-lattice implementation to limit storage growth.

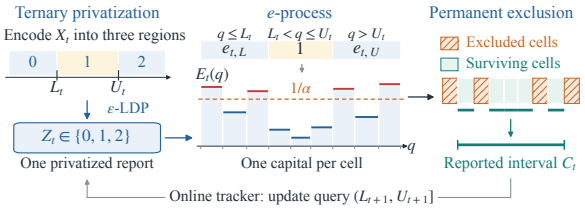

Figure 1: Overview of the proposed method.

Our main contributions are threefold. *(i) Nonasymptotic, anytime-valid LDP quantile inference.* We establish distribution-free simultaneous coverage of at least $1 - \alpha$ under ε-LDP. The guarantee requires neither density regularity nor accurate initialization and remains valid at every almost surely finite data-dependent stopping time. *(ii) Width contraction rate.* Under the acquisition and local density conditions, we establish finite-time acquisition and the asymptotic contraction rate $\mathcal{O}(\sqrt{\log\log t/t})$, with dependence on privacy and the density at the target quantile. *(iii) Uniformly bounded working memory.* The proposed nested-lattice implementation has uniformly bounded working memory. It processes each report once and retains neither raw observations nor report histories. Related work on e-processes is reviewed in Appendix A.

2 Preliminaries

We introduce the definitions of differential privacy (DP), local differential privacy (LDP) and confidence sequences (CSs) to formalize the privacy and coverage requirements of our method.

Definition 1 (DP; Dwork et al., 2006)**.** For $\varepsilon > 0$ and $\delta \in [0, 1)$, a randomized mechanism $\mathcal{M} : \mathcal{X}^n \to \mathcal{O}$ satisfies (ε, δ)-DP if, for every pair of datasets $D, D' \in \mathcal{X}^n$ differing in one individual's

observation and every measurable set $A \subseteq \mathcal{O}$,

$$\Pr\{\mathcal{M}(D) \in A\} \leq e^{\varepsilon} \Pr\{\mathcal{M}(D') \in A\} + \delta,$$

where $\mathcal{O}$ is the output space and the probabilities are taken over the randomization of $\mathcal{M}$. When $\delta = 0$, the mechanism satisfies ε-DP.

Definition 2 (LDP; Xiong et al., 2021)**.** A randomized mechanism $\mathcal{M} : \mathcal{X} \to \mathcal{O}$ satisfies (ε, δ)-LDP if it satisfies (ε, δ)-DP with a single observation as input. The case $\delta = 0$ is termed ε-LDP.

We allow sequential interaction: each user releases one privatized report through a mechanism that may depend on reports from previous users. For any fixed history of these reports, the mechanism must satisfy Definition 2 with respect to the current user's observation (Duchi et al., 2018).

Definition 3 (Nonasymptotic CS; Howard et al., 2021)**.** For $\alpha \in (0, 1)$, we say that the sequence of random intervals $(C_t)_{t \geq 1}$, with $C_t \subseteq \mathbb{R}$, is a $(1 - \alpha)$-CS for a parameter of interest $\theta \in \mathbb{R}$ if

$$\Pr(\theta \in C_t \text{ for all } t \geq 1) \geq 1 - \alpha.$$

Nonasymptotic CSs quantify uncertainty uniformly over time with finite-sample coverage guarantees. In particular, $\Pr(\theta \in C_T) \geq 1 - \alpha$ holds for any almost surely finite stopping time $T \geq 1$, permitting continuous monitoring and data-dependent stopping. In contrast, the validity of asymptotic CSs (Waudby-Smith et al., 2024; Cai et al., 2025) relies on large-sample approximations, which alone do not ensure the same coverage guarantee in finite samples.

3 Methodology

Let $X_1, X_2, \ldots$ be independent and identically distributed observations drawn sequentially from an unknown distribution F. Our goal is to construct a nonasymptotic $(1 - \alpha)$-CS under ε-LDP for the population quantile $q_\tau = \inf\{x \in \mathbb{R} : F(x) \geq \tau\}$, where $\tau \in (0, 1)$. We denote the intervals in this CS by $C_t = [C_{L,t}, C_{U,t}]$, $t \geq 1$, allowing the lower and upper endpoints to equal $-\infty$ and $+\infty$, respectively. We initialize $C_0 = \mathbb{R}$, without requiring a prespecified bounded interval containing q_τ.

Using only privatized reports, our basic idea (Figure 1) is to construct a CS by combining sequential testing with permanent candidate elimination. Specifically, for each quantile candidate $q \in \mathbb{R}$, we consider the following hypothesis test:

$$H_0 : q_\tau = q \quad \text{versus} \quad H_1 : q_\tau \neq q. \tag{1}$$

We implement each sequential test with an e-process $E_t(q)$, permanently excluding q when $E_t(q) \geq 1/\alpha$ for the first time. The interval C_t is then formed as the closed convex hull of the surviving candidates.

A key component of our nonasymptotic CS construction is an online tracker that adaptively selects a query interval $(L_t, U_t]$. This interval serves two purposes: guiding the privatization of X_t and specifying which candidate quantiles are tested for exclusion at each round. In what follows, we first introduce the privatization mechanism and betting-based e-processes in Section 3.1, then describe the online tracker updates in Section 3.2. Finally, we present a nested-lattice algorithm in Section 3.3 to improve computational and memory efficiency.

3.1 Privatization and e-process construction

Let $\mathcal{F}_{t-1}$ represent the information available before round t. Based on this information, the online tracker selects a query interval $(L_t, U_t]$, partitioning $\mathbb{R}$ into $(-\infty, L_t]$, $(L_t, U_t]$, and (U_t, ∞). Under the null in (1), $q \leq L_t$ implies $F(L_t) \geq \tau$, whereas $q > U_t$ implies $F(U_t) < \tau$. We thus consider

$$H^L_{0,t} : F(L_t) \geq \tau \quad \text{versus} \quad H^L_{1,t} : F(L_t) < \tau, \tag{2}$$

$$H^U_{0,t} : F(U_t) \leq \tau \quad \text{versus} \quad H^U_{1,t} : F(U_t) > \tau. \tag{3}$$

Evidence against these nulls supports excluding candidates in $(-\infty, L_t]$ and (U_t, ∞), respectively.

Since $F(L_t) = \mathbb{E}\{I(X_t \leq L_t) \mid \mathcal{F}_{t-1}\}$ and $F(U_t) = \mathbb{E}\{I(X_t \leq U_t) \mid \mathcal{F}_{t-1}\}$, the two one-sided tests (2)-(3) are Bernoulli mean tests. To conduct these tests under ε-LDP, we jointly privatize the two indicators, whose possible outcomes are $(1, 1)$, $(0, 1)$, and $(0, 0)$. This motivates a ternary randomized response mechanism (Waudby-Smith et al., 2023) that provides information for both tests through a single privatized report.

Ternary randomized response. The t-th user encodes X_t as $Y_t = 0$ if $X_t \le L_t$, $Y_t = 1$ if $L_t < X_t \le U_t$, and $Y_t = 2$ if $X_t > U_t$, and reports its randomized version $Z_t \in \{0, 1, 2\}$:

$$\Pr(Z_t = z \mid Y_t = y, \mathcal{F}_{t-1}) = a + \rho I(z = y), \tag{4}$$

where $a = (e^\varepsilon + 2)^{-1}$ and $\rho = (e^\varepsilon - 1)/(e^\varepsilon + 2)$. The mechanism satisfies ε-LDP since its maximum output probability ratio is $(a + \rho)/a = e^\varepsilon$. A detailed derivation is provided in Appendix B.1.

From Z_t, define $B_{t,L} = I(Z_t = 0)$ and $B_{t,U} = I(Z_t \neq 2)$ as privatized counterparts of $I(X_t \le L_t)$ and $I(X_t \le U_t)$, respectively. Their conditional means given $\mathcal{F}_{t-1}$ are $p_{t,L} := a + \rho F(L_t)$ and $p_{t,U} := 2a + \rho F(U_t)$. Thus, with $c_L := a + \rho\tau$ and $c_U := 2a + \rho\tau$, the tests (2)-(3) are equivalently expressed as

$$H_{0,t}^{L,\mathrm{priv}} : p_{t,L} \ge c_L \quad \text{versus} \quad H_{1,t}^{L,\mathrm{priv}} : p_{t,L} < c_L, \tag{5}$$

$$H_{0,t}^{U,\mathrm{priv}} : p_{t,U} \le c_U \quad \text{versus} \quad H_{1,t}^{U,\mathrm{priv}} : p_{t,U} > c_U. \tag{6}$$

Betting-based e-processes. For the private mean tests (5)-(6), we use Bernoulli likelihood-ratio betting factors (Waudby-Smith & Ramdas, 2024):

$$e_{t,j} = \left(\frac{v_{t,j}}{c_j}\right)^{B_{t,j}} \left(\frac{1 - v_{t,j}}{1 - c_j}\right)^{1 - B_{t,j}}, \quad j \in \{L, U\}. \tag{7}$$

Here, $v_{t,j}$ is an $\mathcal{F}_{t-1}$-measurable forecast of $p_{t,j}$ under the alternative. We impose $v_{t,L} \in (0, c_L)$ and $v_{t,U} \in (c_U, 1)$ so that $\mathbb{E}(e_{t,j} \mid \mathcal{F}_{t-1}) \le 1$ under the null, making $e_{t,j}$ a conditional e-value, while $\mathbb{E}(e_{t,j} \mid \mathcal{F}_{t-1}) > 1$ under the alternative supports evidence accumulation.

Our choice of $v_{t,j}$ is guided by the online tracker in Section 3.2, which updates the query interval to bring $\{F(L_t) + F(U_t)\}/2$ closer to τ. For the test (5), this motivates the approximation

$$p_{t,L} = a + \rho F(L_t) \approx a + \rho\tau - \frac{\rho\{F(U_t) - F(L_t)\}}{2} = c_L - \frac{\rho\{F(U_t) - F(L_t)\}}{2}, \tag{8}$$

which indicates that we predict $p_{t,L}$ by estimating $F(U_t) - F(L_t)$. To this end, we use the identity $\mathbb{E}\{I(Z_t = 1) - a \mid \mathcal{F}_{t-1}\} = \rho\{F(U_t) - F(L_t)\}$. Further, we adopt the local linear approximation $\rho\{F(U_t) - F(L_t)\} \approx b(U_t - L_t) = 2br_t$, where $b > 0$ is a slope parameter and $r_t = (U_t - L_t)/2$. Consequently, $\mathbb{E}\{I(Z_t = 1) - a \mid \mathcal{F}_{t-1}\} \approx 2br_t$. Here, we estimate b by constrained least squares using past privatized reports: for $t \ge 2$,

$$\widehat{b}_t = \operatorname*{arg\,min}_{b \ge b_{\min}} \sum_{s=1}^{t-1} \{I(Z_s = 1) - a - 2br_s\}^2 = \max\left\{b_{\min}, \frac{\sum_{s=1}^{t-1} r_s\{I(Z_s = 1) - a\}}{2\sum_{s=1}^{t-1} r_s^2}\right\}, \tag{9}$$

where $b_{\min} > 0$ prevents the e-factors (7) from degenerating to one, and $\widehat{b}_1 = b_0 \ge b_{\min}$. The fitted value $2\widehat{b}_t r_t$ then yields $c_L - \widehat{b}_t r_t$ as a forecast of $p_{t,L}$. Similarly, the test (6) uses $c_U + \widehat{b}_t r_t$.

The distance $|v_{t,j} - c_j|$ controls betting intensity. We moderate the bets by scaling $\widehat{b}_t r_t$ by $\kappa \in (0, 1)$. We then apply an exponential transformation to ensure that the forecasts are valid probabilities satisfying $v_{t,L} \in (0, c_L)$ and $v_{t,U} \in (c_U, 1)$:

$$v_{t,L} = c_L \exp\left(-\frac{\kappa\widehat{b}_t r_t}{c_L}\right), \quad v_{t,U} = 1 - (1 - c_U)\exp\left(-\frac{\kappa\widehat{b}_t r_t}{1 - c_U}\right). \tag{10}$$

These choices approximate $c_L - \kappa\widehat{b}_t r_t$ and $c_U + \kappa\widehat{b}_t r_t$ to first order. It is worth noting that forecast accuracy affects betting efficiency but not the validity of the conditional e-values.

For each candidate q, we select the betting e-factor according to its position relative to the query interval $(L_t, U_t]$:

$$e_t(q) = \begin{cases} e_{t,L}, & q \le L_t, \\ 1, & L_t < q \le U_t, \\ e_{t,U}, & q > U_t. \end{cases} \tag{11}$$

We then combine these factors over time to construct an e-process for q. A direct approach is to use the product $\prod_{s=1}^{t} e_s(q)$. However, this exposes all initial capital to early betting losses, potentially delaying subsequent evidence accumulation. To mitigate this, we retain a prepaid capital reserve and release it gradually. Specifically, letting $R_t = \{1 + \log(t+1)\}^{-1}$ denote the remaining reserve and setting $E_0(q) = R_0 = 1$, we propose the following betting-based e-process:

$$E_t(q) = R_t + \{E_{t-1}(q) - R_t\} e_t(q). \tag{12}$$

The reserve also supports finite-interval acquisition (Theorem C1). Further details and simulations under unfavorable early queries are provided in Appendices B.2 and D.2, respectively.

Nonasymptotic CS construction. At time t, we permanently exclude all candidates q for which $E_s(q) \geq 1/\alpha$ at some round $s \leq t$. The surviving set is

$$S_t = \left\{ q \in \mathbb{R} : \max_{1 \leq s \leq t} E_s(q) < 1/\alpha \right\}. \tag{13}$$

For $S_t \neq \varnothing$, the nonasymptotic CS is given by $C_t = [C_{L,t}, C_{U,t}] = [\inf S_t, \sup S_t]$, with possibly infinite endpoints. The running maximum ensures permanent candidate elimination, so the resulting intervals are nested. If $S_t = \varnothing$, we set $C_u = \varnothing$ for $u \geq t$ and terminate the procedure.

3.2 Online tracking and query updates

This section describes how the query interval $(L_t, U_t]$ is updated adaptively in two stages. Define the first finite-interval time as

$$\tilde{\sigma} = \inf\{t \geq 1 : C_t \text{ is bounded and nonempty}\}, \quad \inf \varnothing = \infty. \tag{14}$$

Stage 1: finite CS acquisition. Choose an arbitrary initial position $\theta_1 \in \mathbb{R}$ and a fixed radius $h_1 > 0$. For $t \leq \tilde{\sigma}$, we use the proposed center $\mu_t = \theta_t$ and radius $h_t = h_1$, updating only the center. Our objective is to bring $\{F(L_t) + F(U_t)\}/2$ closer to τ, consistent with the approximation in (8).

To determine the update direction from privatized reports, let $K_t := (B_{t,L} + B_{t,U})/2 = 1 - Z_t/2$ and $c_0 = (c_L + c_U)/2$. Then $\mathbb{E}(K_t - c_0 \mid \mathcal{F}_{t-1}) = \rho[\{F(L_t) + F(U_t)\}/2 - \tau]$, indicating that the query should move left when this mean is positive and right when it is negative. Accordingly, after receiving the privatized report Z_t, we use the following Newton-type stochastic update:

$$\theta_{t+1} = \theta_t - \frac{K_t - c_0}{(t + t_0)\widehat{b}_t}, \tag{15}$$

which defines the online tracker in Stage 1. Here, $\widehat{b}_t$ is the local slope estimate defined in (9), while the fixed integer $t_0 \geq 0$ controls the initial step size. For $t \leq t_0$, we set $e_{t,L} = e_{t,U} = 1$ while continuing the tracker and slope updates.

Stage 2: finite CS contraction. For $t > \tilde{\sigma}$, the online tracker in Stage 2 uses the midpoint of C_{t-1} as the proposed query center and a fixed fraction of its width as the query width: $\mu_t = (C_{L,t-1} + C_{U,t-1})/2$ and $h_t = \lambda(C_{U,t-1} - C_{L,t-1})/2$. The parameter $\lambda \in (0, 1)$ specifies the ratio of the proposed query width to the CS width, and we set $\lambda = e^{-1/2}$. As the CS contracts, the proposed query interval shrinks accordingly, allowing previously neutral candidates to receive betting updates.

Detailed update rules and further discussion of the choices of h_1 and λ are provided in Appendix B.3. The proposed query endpoints are $L_t = \mu_t - h_t$ and $U_t = \mu_t + h_t$, which are rounded in Section 3.3 for the lattice implementation.

3.3 Nested-lattice implementation

As the query endpoints L_t and U_t are updated over time, they progressively divide $\mathbb{R}$ into smaller intervals, which we call cells. Candidates within each cell share the same betting-factor history and therefore the same capital. Figure 2 illustrates how evidence is updated within each cell. Let M_t denote the number of surviving cells at time t, namely those not yet permanently excluded. Storing the endpoints and one shared capital value for each surviving cell requires $\mathcal{O}(M_t)$ memory. As new query endpoints split existing cells, this storage requirement can grow over time.

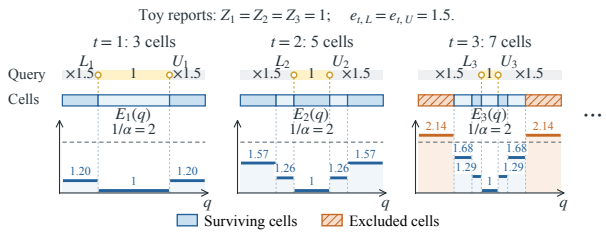

Figure 2: Cellwise evidence updates in a three-round toy example with threshold $1/\alpha = 2$. At $t = 3$, two cells are excluded and five survive.

To limit this growth, we round query endpoints outward onto nested lattices, allowing nearby endpoints to share grid points and reducing the creation of new cells. Specifically, starting from $\Delta_0 = \infty$, we set

$$\Delta_t = \min\left\{\Delta_{t-1}, 2^{\lfloor \log_2(\alpha h_t)\rfloor}\right\}, \quad \widetilde{L}_t = \Delta_t \left\lfloor \frac{\mu_t - h_t}{\Delta_t} \right\rfloor, \quad \widetilde{U}_t = \Delta_t \left\lceil \frac{\mu_t + h_t}{\Delta_t} \right\rceil. \tag{16}$$

Here, Δ_t controls the grid spacing. A larger Δ_t limits storage growth, whereas a smaller Δ_t provides finer query resolution. We decrease Δ_t as h_t shrinks to balance storage efficiency and query precision. Later, we show in Theorem C2 that this nested-lattice implementation has memory bounded uniformly over time.

In the lattice implementation, $\widetilde{L}_t$ and $\widetilde{U}_t$ replace L_t and U_t throughout the procedure. The actual radius $r_t = (\widetilde{U}_t - \widetilde{L}_t)/2$ is used for slope estimation in (9) and betting in (10). We provide the full algorithm and implementation details in Appendix B.4.

4 Theoretical Guarantees

We establish ε-LDP and distribution-free, nonasymptotic simultaneous coverage for the proposed method, and derive its width contraction rate under suitable regularity conditions. Theoretical results on finite-interval acquisition and uniformly bounded working memory are provided in Appendix C.

Theorem 1 (ε-LDP and nonasymptotic anytime-valid coverage)**.** *Fix $\varepsilon > 0$ and $\alpha, \tau \in (0,1)$. The proposed algorithm satisfies ε-LDP for each user, and the released CS preserves this guarantee by post-processing. For every distribution F on $\mathbb{R}$, we have*

$$\Pr\left(q_\tau \in C_t \text{ for all } t \geq 1\right) \geq \Pr\left(q_\tau \in S_t \text{ for all } t \geq 1\right) \geq 1 - \alpha,$$

where S_t is the surviving set in (13) and $C_t = [\inf S_t, \sup S_t]$ when $S_t \neq \varnothing$, with $C_t = \varnothing$ otherwise. Consequently, $\Pr(q_\tau \in C_T) \geq 1 - \alpha$ for every almost surely finite stopping time $T \geq 1$.

Remark 1. Theorem 1 establishes a nonasymptotic $(1-\alpha)$-CS for q_τ under ε-LDP. Coverage follows from the supermartingale property of $E_t(q_\tau)$, without requiring accurate initialization or consistent slope estimation. Compared to the asymptotic LDP quantile CSs of Cai et al. (2025), our guarantee requires no density regularity and remains valid at finite sample sizes under data-dependent stopping, including for discrete distributions and discontinuous densities. In the non-private case with $a = 0$ and $\rho = 1$, our method offers a betting-based alternative to the order-statistics-based CS of Howard & Ramdas (2022), retaining distribution-free, nonasymptotic coverage.

We next characterize how rapidly the CS contracts around q_τ. Unlike the coverage guarantee, this result requires local regularity of F.

Assumption 1 (Local regularity)**.** *The distribution F has a continuous positive density f in a neighborhood of q_τ, and $0 < b_{\min} \leq \rho f(q_\tau)$, where $b_{\min}$ is the lower bound imposed on $\widehat{b}_t$ in (9).*

The density condition ensures that $F(x) - \tau \approx f(q_\tau)(x - q_\tau)$ near q_τ. Since $f(q_\tau) > 0$, bounds on $|F(x) - \tau|$ translate into bounds on $|x - q_\tau|$, allowing us to control the interval width. The condition on $b_{\min}$ concerns slope estimation: a lower bound above the local private slope $\rho f(q_\tau)$ prevents consistent estimation in (9) and may lead to overly aggressive betting. Assumption 1 is used in the contraction-rate proof, but is not required for the coverage guarantee.

Theorem 2. *Suppose Assumption 1 and the acquisition condition in Assumption C1 of Appendix C.1 hold. Let $\alpha \in (0,1)$ and $\lambda \in (1/2, 1/(1+\alpha))$ be fixed. On the event $\{q_\tau \in S_t \text{ for all } t \geq 1\}$, the CS produced by the nested-lattice implementation is eventually bounded and its full width $W_t = C_{U,t} - C_{L,t}$ satisfies, almost surely,*

$$W_t = \mathcal{O}\left(\frac{1}{\rho f(q_\tau)}\sqrt{\frac{\log\log t}{t}}\right), \quad t \to \infty.$$

Remark 2. Theorem 2 shows that the proposed LDP quantile CS contracts at the iterated-logarithm rate $\mathcal{O}(\sqrt{\log\log t/t})$. Under the local density condition, this has the same order as the non-private quantile CS of Howard & Ramdas (2022) and the asymptotic stitched LDP quantile CS of Cai et al. (2025). The privacy budget ε and the local density $f(q_\tau)$ affect only the constant, not the contraction rate in t.

In addition, Theorem C1 in Appendix C.1 guarantees finite-time acquisition of a bounded, nonempty CS with probability at least $1-\alpha$ under the acquisition condition, without requiring a prespecified interval containing q_τ. Theorem C2 in Appendix C.2 shows that the working memory of the nested-lattice algorithm remains bounded as the number of users increases. In contrast, the standard exact order-statistic implementation of Howard & Ramdas (2022) requires $\mathcal{O}(T)$ memory for T users.

5 Numerical Experiments

We evaluate the proposed method through simulations and a real data analysis, examining its finite-sample performance and advantages relative to asymptotic LDP confidence sequences (CSs).

Basic setups. The simulation data $X_1, \ldots, X_T$ are sequentially and independently generated from a common distribution F, considering the following three cases: (i) the standard normal distribution $\mathcal{N}(0,1)$; (ii) the standard Cauchy distribution $\mathcal{C}(0,1)$; (iii) the lognormal distribution $\mathrm{Logn}(0,1)$. Following Cai et al. (2025), we set the total number of users to $T = 5 \times 10^6$. Moreover, we consider three quantile levels $\tau \in \{0.3, 0.5, 0.7\}$ and four privacy budgets $\varepsilon \in \{0.5, 1, 2, 3\}$, where a smaller value of ε imposes stronger privacy protection. All CSs are constructed at the 95% nominal level.

We conduct $L = 1{,}000$ independent replications and evaluate the constructed CSs using the empirical simultaneous coverage probability (SCP) and the average interval width (AW):

$$\mathrm{SCP}(t) = \frac{1}{L}\sum_{l=1}^{L} I\left(q_\tau \in C_s^{(l)} \text{ for all } 1 \leq s \leq t\right), \quad \mathrm{AW}(t) = \frac{1}{L}\sum_{l=1}^{L} W_t^{(l)},$$

where $C_t^{(l)}$ denotes the CS after t users in the l-th replication, and $W_t^{(l)}$ is its full width. SCP and AW are reported only at checkpoints where every replication has a bounded interval; before that, AW is recorded as infinite.

Comparison methods. We primarily compare our method with the Gaussian-mixture asymptotic CS of Cai et al. (2025) (Cai GM), motivated by its favorable finite-sample performance relative to their stitched construction (Cai stitched). The baseline initial position is 0 for our tracker and all Cai chains. For our method, we set $t_0 = 64$, $b_0 = 0.18$, $b_{\min} = 0.01$, $\kappa = 0.75$, and $\lambda = e^{-1/2}$. For Cai GM, we adopt their increasing-chain construction, step sizes $\eta_j = j^{-0.6}$, and prescribed burn-in rule, with boundary parameter $\rho_{\mathrm{GM}} = 0.001$ in simulations and 0.01 in the real data analysis. Additional comparisons with Cai stitched are provided in Appendix D.1. It is worth noting that we record infinite interval width before our method obtains a finite interval and during the burn-in period of Cai GM and Cai stitched.

Experiments. The main text presents two simulation experiments. *Experiment 1* examines simultaneous coverage and CS width of our method under $\mathcal{N}(0,1)$ across privacy budgets. To our knowledge, no existing nonasymptotic LDP quantile CS is available for comparison. We therefore include the non-private nonasymptotic quantile CS of Howard & Ramdas (2022) as a benchmark. The results are summarized in Figure 3 and Table D1 in Appendix D. *Experiment 2* compares our method with Cai GM. First, we examine robustness to a density discontinuity using a piecewise-scaled normal distribution: $X = Z$ for $Z < 0$ and $X = 2Z$ otherwise, where $Z \sim \mathcal{N}(0,1)$, so the density

jumps at the median $q_{0.5} = 0$ with $f(0^-) = \phi(0)$ and $f(0^+) = \phi(0)/2$, where ϕ denotes the standard normal density. Second, we assess initialization sensitivity under $\mathcal{N}(0, 1)$ at $\tau = 0.5$, varying the initial position over $\{0, 4, 8, 12, 16\}$ to progressively increase its distance from the true $q_{0.5} = 0$. The results are reported in Figures 4-5.

For a comprehensive evaluation, we supplement *Experiment 1* with results for the standard Cauchy, lognormal, and three-component Beta mixture distributions, together with a scale-sensitivity analysis under the normal distribution in Appendix D.1. We also conduct five additional experiments. *Experiment A1* evaluates the role of prepaid capital reserves by comparing our method with and without reserves (Figure D7). *Experiment A2* examines the effect of lattice rounding on CS paths (Figure D8). *Experiment A3* tracks the number of surviving cells across privacy budgets to assess memory requirements (Figure D9). *Experiment A4* examines whether CS widths shrink at a rate consistent with our theory (Figures D10-D11). *Experiment A5* examines sensitivity to tuning parameters (Figures D12-D14).

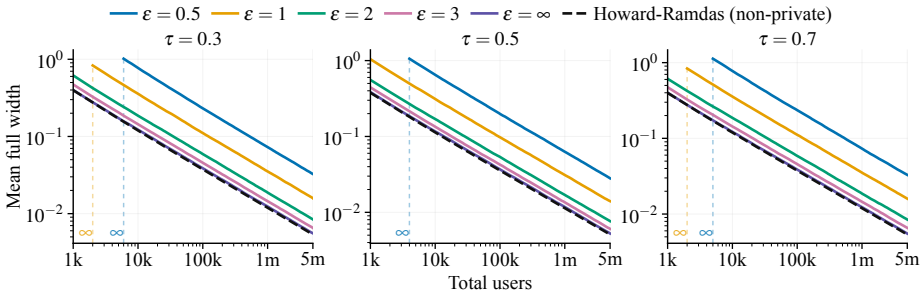

Figure 3: AW under $\mathcal{N}(0, 1)$ across privacy budgets, with the non-private Howard–Ramdas benchmark. Here $\varepsilon = \infty$ denotes the non-private case; k and m denote 10^3 and 10^6 users; both axes are logarithmic.

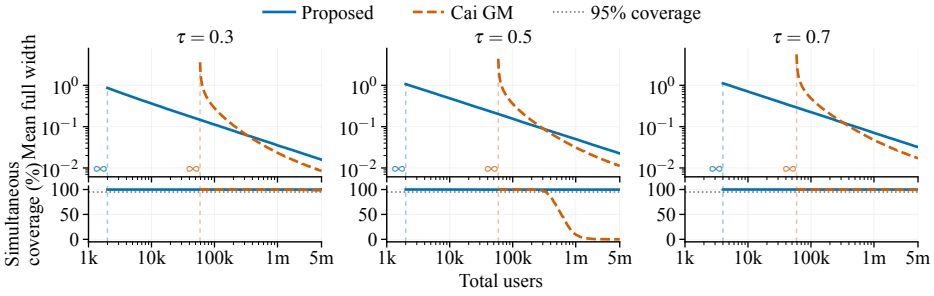

Figure 4: Piecewise-scaled normal data at $\varepsilon = 1$: AW (top) and SCP (bottom) for $\tau = 0.3, 0.5, 0.7$. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

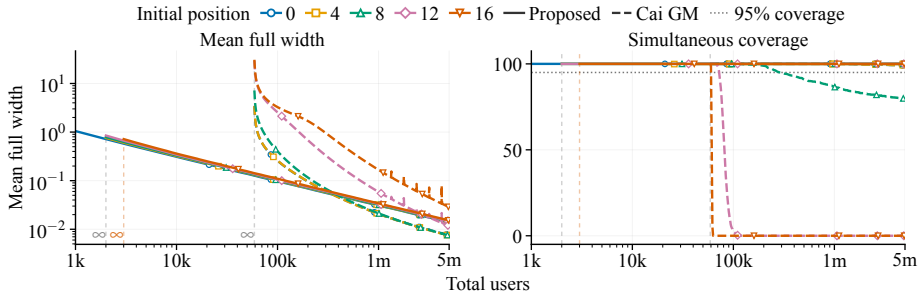

Figure 5: Initialization sensitivity under $\mathcal{N}(0, 1)$ at $\tau = 0.5$ and $\varepsilon = 1$: AW (left) and SCP (right). Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

Results. In *Experiment 1*, Figure 3 shows that AW decreases as the sample size increases, with larger privacy budgets yielding narrower intervals. In the non-private setting, the AW of our method is comparable to that of the Howard–Ramdas benchmark. Unlike its exact order-statistic implementation, which uses $\mathcal{O}(t)$ memory, our construction requires only bounded working memory (Theorem C2). In addition, Table D1 indicates that the empirical SCP remains above the nominal level across all quantile levels and privacy budgets, consistent with Theorem 1.

In *Experiment 2*, Figure 4 shows that our method obtains finite intervals much earlier than Cai GM, which requires a substantial burn-in period under its prescribed implementation. Our method also achieves smaller AW during the early stages while maintaining SCP above the nominal level. When the density is discontinuous at the median, Cai GM exhibits severe undercoverage, whereas our method retains coverage without requiring density continuity. Moreover, Figure 5 indicates that our method shows stable performance across initial positions, consistent with the adaptive updates of our tracker in Section 3.2. In contrast, Cai GM is sensitive to distant initialization: starting from 12 or 16 leads to undercoverage and wider AW. These results demonstrate the practical benefits of our nonasymptotic construction when inference is needed before initialization effects have dissipated.

Experiments A1-A5 highlight the benefits of prepaid capital reserves and lattice rounding, support consistency with our theoretical results, and demonstrate robustness to tuning parameter choices. Detailed results and discussion are provided in Appendix D.

Real data analysis. We analyze annual wage income using the 2018 American Community Survey California public-use microdata (U.S. Census Bureau, 2018). Individual earnings reveal personal financial information, motivating privacy protection for statistical analysis. We include individuals aged at least 18 with positive wage income, yielding 179,799 records, and adjust incomes to 2018 dollars. We perform quantile inference on $X_t = \log(\text{wage}_t/50{,}000)$ and report CS bounds on the original dollar scale. The empirical median \$40,523.88 serves only as a descriptive reference.

We compare our method with Cai GM at $\tau = 0.5$ and $\varepsilon = 1$. Figure 6 shows that our method obtains a finite interval after 348 users, compared with 2,125 for Cai GM under its prescribed burn-in. Both methods subsequently produce shrinking intervals around the empirical median. Although Cai GM is slightly narrower at the end, our method provides informative inference substantially earlier, highlighting the finite-sample advantage of our nonasymptotic approach.

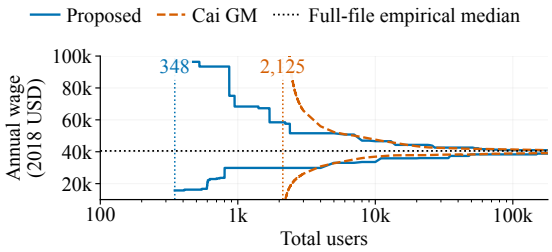

Figure 6: CS trajectories for the wage at $\tau = 0.5$ and $\varepsilon = 1$. The dotted horizontal line marks the full-data empirical median. Here k denotes 10^3 users; user count uses a log scale.

6 Conclusions and Discussions

This work addresses the need for nonasymptotic, time-uniform quantile inference under LDP. Our framework connects locally private quantile inference with betting e-processes through adaptive queries: instead of privatizing an existing non-private method, it builds the CS solely from privatized reports by sequentially testing candidate quantiles. The resulting CS is valid for every distribution and at every almost surely finite stopping time, and the nested-lattice implementation requires only bounded memory. Under the stated conditions, the width contracts at the iterated-logarithm rate, the same rate attained by non-private quantile CSs, so privatization affects only the constant. Although the CS remains valid in finite samples, where asymptotic methods may undercover, it is conservative and can be wider than asymptotic alternatives at later stages. Future work includes scale-invariant tuning, improving betting efficiency, and extending the framework to multiple quantiles with shared privatized reports.

AI USE STATEMENT

Generative AI tools were used to polish the manuscript, identify relevant literature, assist in developing and refining mathematical proofs, review and debug code, and check experimental results. The authors independently reviewed all AI-assisted text and code, verified the references against the original sources, checked the mathematical arguments, and confirmed the reported results against the experimental outputs. The authors take full responsibility for the entire content of this work, including its mathematical claims, proofs, code, and experimental results.

REFERENCES

Leheng Cai, Qirui Hu, Juntao Sun, and Shuyuan Wu. Time-uniform and asymptotic confidence sequence of quantile under local differential privacy. In *Advances in Neural Information Processing Systems*, volume 38, pp. 114488–114520, 2025.

Leheng Cai, Qirui Hu, and Shuyuan Wu. Federated learning of quantile inference under local differential privacy. In *The Fourteenth International Conference on Learning Representations*, 2026.

Likai Chen, Georg Keilbar, and Wei Biao Wu. Recursive quantile estimation: Non-asymptotic confidence bounds. *Journal of Machine Learning Research*, 24(91):1–25, 2023.

D. A. Darling and Herbert Robbins. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences*, 58(1):66–68, 1967.

Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–201, 2018.

Rick Durrett. *Probability: Theory and Examples*. Cambridge University Press, 2019.

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

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, pp. 265–284, 2006.

Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1054–1067, 2014.

Peter Grünwald, Rianne de Heide, and Wouter Koolen. Safe testing. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(5):1091–1128, 2024.

Steven R. Howard and Aaditya Ramdas. Sequential estimation of quantiles with applications to A/B testing and best-arm identification. *Bernoulli*, 28(3):1704–1728, 2022.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.

Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics*, 49(2):1055–1080, 2021.

Ramesh Johari, Pete Koomen, Leonid Pekelis, and David Walsh. Always valid inference: Continuous monitoring of A/B tests. *Operations Research*, 70(3):1806–1821, 2022.

Roger Koenker. *Quantile Regression*. Cambridge University Press, 2005.

Yi Liu, Qirui Hu, Lei Ding, and Linglong Kong. Online local differential private quantile inference via self-normalization. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 21698–21714, 2023.

Yi Liu, Qirui Hu, and Linglong Kong. Tuning-free estimation and inference of cumulative distribution function under local differential privacy. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 31147–31164, 2024.

Yan Michalevsky, Aaron Schulman, Gunaa Arumugam Veerapandian, Dan Boneh, and Gabi Nakibly. PowerSpy: Location tracking using mobile device power analysis. In *24th USENIX Security Symposium (USENIX Security 15)*, pp. 785–800, 2015.

Aaditya Ramdas and Ruodu Wang. Hypothesis testing with e-values. *Foundations and Trends in Statistics*, 1(1–2):1–390, 2025.

Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.

Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 184(2):407–431, 2021.

Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, Bayes factors and p-values. *Statistical Science*, 26(1):84–101, 2011.

Jaehyeok Shin, Aaditya Ramdas, and Alessandro Rinaldo. E-detectors: A nonparametric framework for sequential change detection. *The New England Journal of Statistics in Data Science*, 2(2): 229–260, 2024.

Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. Privacy loss in Apple's implementation of differential privacy on MacOS 10.12. *arXiv preprint arXiv:1709.02753*, 2017.

U.S. Census Bureau. 2018 ACS 1-year PUMS: California person file and data dictionary, 2018. URL https://www.census.gov/programs-surveys/acs/microdata/documentation.2018.html.

Jean Ville. *Étude Critique de la Notion de Collectif*. Gauthier-Villars, Paris, 1939.

Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *Annals of Statistics*, 49(3):1736–1754, 2021.

Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1):1–27, 2024.

Ian Waudby-Smith, Steven Wu, and Aaditya Ramdas. Nonparametric extensions of randomized response for private confidence sets. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pp. 36748–36789, 2023.

Ian Waudby-Smith, David Arbour, Ritwik Sinha, Edward H. Kennedy, and Aaditya Ramdas. Time-uniform central limit theory and asymptotic confidence sequences. *Annals of Statistics*, 52(6): 2613–2640, 2024.

Xingxing Xiong, Shubo Liu, Dan Li, Zhaohui Cai, and Xiaoguang Niu. A comprehensive survey on local differential privacy. *Security and Communication Networks*, 2021(1), 2021.

A Related work on e-processes

e-values and anytime-valid inference. An e-value is a nonnegative statistic whose expectation is at most one under the null hypothesis, and an e-process is a nonnegative process whose value at every stopping time is an e-value (Vovk & Wang, 2021; Grünwald et al., 2024; Ramdas & Wang, 2025). The canonical example is a test martingale, a nonnegative martingale under the null that starts at one (Shafer et al., 2011). Ville's inequality (Ville, 1939) bounds the probability that a nonnegative supermartingale starting at one ever exceeds $1/\alpha$ by α, so thresholding an e-process yields a test that remains valid under continuous monitoring and optional stopping (Ramdas et al., 2023). Inverting such tests yields CSs (Howard et al., 2020; 2021). Our coverage argument follows this route: $E_t(q_\tau)$ is a nonnegative supermartingale, and the CS is the closed convex hull of the candidates whose e-processes have never reached $1/\alpha$ (Theorem 1).

Testing by betting. Shafer (2021) interprets an e-process as the wealth of a gambler betting against the null at fair odds, so large wealth is evidence against it. Waudby-Smith & Ramdas (2024) develop predictable betting strategies that yield tight, variance-adaptive CSs for bounded means, and Waudby-Smith et al. (2023) extend betting CSs to locally privatized bounded data. Our Bernoulli likelihood-ratio factors follow the same principle but differ in two respects. First, the one-sided nulls move with the adaptively chosen query, and each candidate outside the query is tested through the side on which it lies. Second, each ternary report updates the e-processes of all candidates. Forecast accuracy affects betting efficiency but not coverage validity.

Combining e-processes. Convex combinations of e-processes remain e-processes, so evidence can be pooled across starting times. The e-detectors of Shin et al. (2024) also combine e-processes started at different times, but target the average run length in change detection. Our reserve instead allocates a unit budget across starting times (B1), which keeps $E_t(q_\tau)$ an e-process and shields later allocations from early betting losses.

B Further details of the proposed method

B.1 Privacy guarantee of ternary randomized response

Recall that the mechanism in (4) releases $Z_t \in \{0,1,2\}$ according to

$$\Pr(Z_t = z \mid Y_t = y, \mathcal{F}_{t-1}) = a + \rho I(z = y) = \begin{cases} \dfrac{e^\varepsilon}{e^\varepsilon + 2}, & z = y, \\ \dfrac{1}{e^\varepsilon + 2}, & z \neq y, \end{cases}$$

where $a = (e^\varepsilon + 2)^{-1}$ and $\rho = (e^\varepsilon - 1)/(e^\varepsilon + 2)$. We apply the likelihood-ratio argument for randomized response (Waudby-Smith et al., 2023, Appendix A.2).

For any fixed history $\mathcal{F}_{t-1}$, the query interval is fixed and Y_t is determined by X_t. Thus, for any $x, x' \in \mathbb{R}$ and $z \in \{0,1,2\}$, we have

$$\frac{1}{e^\varepsilon + 2} \le \Pr(Z_t = z \mid X_t = x, \mathcal{F}_{t-1}) \le \frac{e^\varepsilon}{e^\varepsilon + 2},$$
$$\frac{\Pr(Z_t = z \mid X_t = x, \mathcal{F}_{t-1})}{\Pr(Z_t = z \mid X_t = x', \mathcal{F}_{t-1})} \le \frac{e^\varepsilon/(e^\varepsilon + 2)}{1/(e^\varepsilon + 2)} = e^\varepsilon.$$

Consequently, for any $A \subseteq \{0,1,2\}$,

$$\begin{aligned} \Pr(Z_t \in A \mid X_t = x, \mathcal{F}_{t-1}) &= \sum_{z \in A} \Pr(Z_t = z \mid X_t = x, \mathcal{F}_{t-1}) \\ &\le e^\varepsilon \sum_{z \in A} \Pr(Z_t = z \mid X_t = x', \mathcal{F}_{t-1}) \\ &= e^\varepsilon \Pr(Z_t \in A \mid X_t = x', \mathcal{F}_{t-1}). \end{aligned}$$

Therefore, the mechanism satisfies ε-LDP conditional on every history $\mathcal{F}_{t-1}$.

B.2 PREPAID CAPITAL RESERVES

In Section 3.1, we introduced the e-process (12) with a prepaid capital reserve. Here, we examine how its capital allocation supports recovery from early betting losses. Define

$$R_t = \frac{1}{1+\log(t+1)}, \quad \beta_s = R_{s-1} - R_s > 0,$$

where β_s is the capital released at round s. These allocations satisfy

$$R_t + \sum_{s=1}^{t} \beta_s = 1, \quad \sum_{s=1}^{\infty} \beta_s = 1.$$

Thus, all allocations come from the initial unit capital, with no additional funding. Writing (12) as $E_t(q) - R_t = \{E_{t-1}(q) - R_{t-1} + \beta_t\} e_t(q)$ and iterating from $E_0(q) = R_0 = 1$ gives

$$E_t(q) = R_t + \sum_{s=1}^{t} \beta_s \prod_{u=s}^{t} e_u(q), \tag{B1}$$

which shows that each allocation β_s participates only in bets at rounds $u \geq s$ and is thus unaffected by betting losses before round s.

More precisely, for every $1 \leq s \leq t$,

$$\log E_t(q) \geq \sum_{u=s}^{t} \log e_u(q) - \log(1/\beta_s).$$

Consequently, $\sum_{u=s}^{t} \log e_u(q) \geq \log(1/\alpha) + \log(1/\beta_s)$ is sufficient for rejection, regardless of betting losses before round s. The additional term $\log(1/\beta_s)$ reflects the cost of starting with capital β_s rather than one. For the chosen reserve,

$$\beta_s = \frac{\log(1+1/s)}{(1+\log s)\{1+\log(s+1)\}} \sim \frac{1}{s(\log s)^2},$$
$$\log(1/\beta_s) = \log s + 2\log\log s + o(1), \quad s \to \infty.$$

This provides a mechanism for recovery after unfavorable early queries. The trade-off is that withholding capital can slow evidence accumulation when early bets are already informative. The simulations in Appendix D.2 examine these effects.

B.3 ONLINE TRACKER UPDATES AND PARAMETER CHOICES

Initial query radius. Recall that the tracker uses the average of the two private indicators derived from the same report:

$$K_t = \frac{B_{t,L} + B_{t,U}}{2} = 1 - \frac{Z_t}{2}.$$

With $c_L = a + \rho\tau$, $c_U = 2a + \rho\tau$, and $c_0 = (c_L + c_U)/2$, the centered signal $K_t - c_0$ guides the movement of the query center. To select the initial radius, we first quantify the variability of this signal under a reference configuration.

The randomized response mechanism gives

$$\Pr(Z_t = z \mid \mathcal{F}_{t-1}) = a + \rho \Pr(Y_t = z \mid \mathcal{F}_{t-1}), \quad z \in \{0,1,2\}.$$

For calibration, consider the reference limit $F(L_t) = F(U_t) = \tau$. Since Y_t records which query region contains X_t, we obtain

$$\Pr(Z_t = 0 \mid \mathcal{F}_{t-1}) = a + \rho F(L_t) = a + \rho\tau = c_L,$$
$$\Pr(Z_t = 1 \mid \mathcal{F}_{t-1}) = a + \rho\{F(U_t) - F(L_t)\} = a,$$
$$\Pr(Z_t = 2 \mid \mathcal{F}_{t-1}) = a + \rho\{1 - F(U_t)\} = a + \rho(1-\tau) = 1 - c_L - a,$$

where the last equality uses $3a + \rho = 1$.

The signal K_t takes values 1, $1/2$, and 0 when $Z_t = 0$, 1, and 2, respectively. Its conditional moments under this reference configuration are therefore

$$\mathbb{E}(K_t \mid \mathcal{F}_{t-1}) = c_L + \frac{a}{2} = c_0,$$
$$\mathbb{E}(K_t^2 \mid \mathcal{F}_{t-1}) = c_L + \frac{a}{4}.$$

Hence, the reference variance of the centered update signal is

$$V_\tau := \mathrm{Var}(K_t - c_0 \mid \mathcal{F}_{t-1}) = c_L + \frac{a}{4} - c_0^2 = c_0(1-c_0) - \frac{a}{4}.$$

Under the tracker's local linear approximation, b_0 represents the initial change in the mean private signal per unit change in location. Dividing the reference standard deviation $\sqrt{V_\tau}$ by b_0 therefore converts signal fluctuations into a corresponding location scale. We choose the initial radius as

$$h_1 = \frac{2\sqrt{2}}{\sqrt{1+t_0}} \frac{\sqrt{V_\tau}}{b_0}. \tag{B2}$$

Thus, greater signal variability or a smaller initial slope leads to a wider query, while a larger offset t_0 reduces its initial width. The multiplier $2\sqrt{2}$ is fixed for calibration. This rule determines the initial query width without requiring the query interval to contain q_τ. The reference condition serves only to calibrate h_1; it is not assumed to describe the actual queries. Notably, the finite-interval acquisition guarantee for the chosen radius is established under Assumption C1.

Update order and stage transition. At round t, we first compute $\widehat{b}_t$ using the information in $\mathcal{F}_{t-1}$. The proposed center and radius are

$$(\mu_t, h_t) = \begin{cases} (\theta_t, h_1), & C_{t-1} \text{ is unbounded}, \\ \left(\dfrac{C_{L,t-1} + C_{U,t-1}}{2}, \dfrac{\lambda(C_{U,t-1} - C_{L,t-1})}{2} \right), & C_{t-1} \text{ is bounded and nonempty}. \end{cases}$$

For the lattice implementation, we round the proposed endpoints as in (16) and compute the actual radius r_t before forming the forecasts $v_{t,L}$ and $v_{t,U}$ for the private Bernoulli means in (10). Thus, neither the query nor the forecasts use the current report Z_t.

Figure B1 summarizes the update cycle. Each report Z_t updates the candidate capitals by (12) and yields C_t through (13). In Stage 1, it also updates θ_{t+1} by (15), using $\widehat{b}_t$ based on $\mathcal{F}_{t-1}$. After the first finite CS is obtained, Stage 2 begins at $\tilde{\sigma}+1$ with a new query determined by $C_{\tilde{\sigma}}$. Subsequent queries are updated only when a CS endpoint changes; otherwise, the query and lattice remain fixed. Slope estimation and evidence accumulation continue at every round, with no reset between stages.

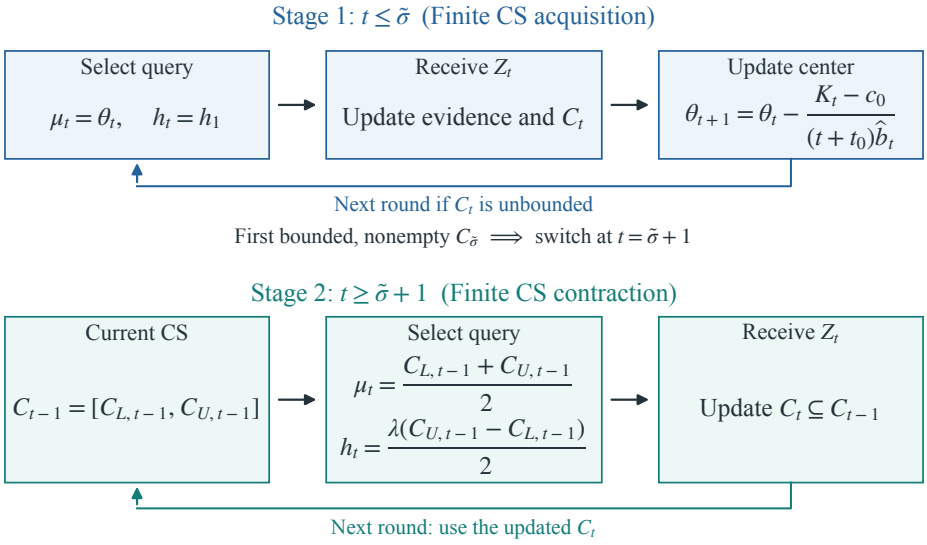

Figure B1: Illustration of the online tracker in Section 3.2.

Choice of the query fraction. We choose λ to balance evidence accumulation and interval contraction: a wider query can provide stronger evidence against outer candidates, whereas a narrower query allows greater contraction once those candidates are excluded.

Let ω denote the current CS half-width, and set the proposed query radius to $h = \lambda\omega$. To guide the choice of λ, consider the centered query $(L, U] = (q_\tau - h, q_\tau + h]$. For this calculation, assume that F has a density f that is continuous and positive at q_τ. Since $F(q_\tau) = \tau$, as $h \to 0$,

$$F(L) - \tau = -\int_{q_\tau - h}^{q_\tau} f(x)\,dx = -f(q_\tau)h + o(h),$$
$$F(U) - \tau = \int_{q_\tau}^{q_\tau + h} f(x)\,dx = f(q_\tau)h + o(h).$$

The private means therefore satisfy

$$p_L - c_L = \rho\{F(L) - \tau\} = -\rho f(q_\tau)h + o(h),$$
$$p_U - c_U = \rho\{F(U) - \tau\} = \rho f(q_\tau)h + o(h).$$

If the slope forecast satisfies $\widehat{b} = \rho f(q_\tau) + o(1)$, a second-order expansion of the expected log betting factor gives

$$\mathbb{E}\log e_j = \frac{\kappa(2-\kappa)\rho^2 f(q_\tau)^2}{2c_j(1-c_j)}h^2 + o(h^2), \quad j \in \{L, U\}.$$

Thus, the leading evidence-growth rate is proportional to h^2.

If candidates outside the query are excluded, the CS half-width is reduced from ω to at most h. We measure this potential contraction by $\log(\omega/h)$ and combine it with the leading evidence-growth rate to define the criterion

$$J(\lambda) = \frac{h^2}{\omega^2}\log\frac{\omega}{h} = \lambda^2\log(1/\lambda).$$

Differentiating gives

$$J'(\lambda) = \lambda\{2\log(1/\lambda) - 1\},$$

yielding

$$\lambda = \operatorname*{arg\,max}_{0<u<1} J(u) = e^{-1/2}.$$

We therefore use $\lambda = e^{-1/2}$ as the default query fraction. Its optimality applies to this local criterion.

B.4 Full algorithm and implementation

Algorithm B1 summarizes the proposed nested-lattice procedure, in which the rounded endpoints $\widetilde{L}_t$ and $\widetilde{U}_t$ replace L_t and U_t.

Algorithm B1 Nonasymptotic quantile CS under ε-LDP

Require: $\tau, \varepsilon, \alpha$; initial position θ_1; parameters $b_0, b_{\min}, t_0, \kappa, \lambda$.

Compute a, ρ, c_L, c_U as in Section 3.1, h_1 by (B2), and $c_0 = (c_L + c_U)/2$.
Initialize one surviving cell $\mathbb{R}$ with total capital $E_0(q) = 1$, $C_0 = \mathbb{R}$, $R_0 = 1$, and $\Delta_0 = \infty$.
Initialize the two cumulative sums used to compute $\widehat{b}_t$ in (9) to zero.
for $t = 1, 2, \ldots$ **do**
    Compute $\widehat{b}_t$ by (9), with $\widehat{b}_1 = b_0$.
    Select or retain the query as illustrated in Figure B1, using the lattice rounding rule (16).
    Set $r_t = (\widetilde{U}_t - \widetilde{L}_t)/2$ and compute $v_{t,L}, v_{t,U}$ by (10).
    Receive the user's privatized report Z_t through (4).
    Split surviving cells at $\widetilde{L}_t, \widetilde{U}_t$ where needed; each child inherits its parent's capital.
    if $t \le t_0$ **then**
        Set $e_{t,L} = e_{t,U} = 1$.
    else
        Compute the e-factors $e_{t,L}, e_{t,U}$ by (7).
    end if
    Update the remaining reserve $R_t = \{1 + \log(t+1)\}^{-1}$.
    Assign the e-factors by (11) and update each cell's capital by (12).
    Permanently delete every cell whose updated capital is at least $1/\alpha$.
    if no cells remain **then**
        Set $C_u = \varnothing$ for all $u \ge t$ and terminate.
    end if
    Let S_t be the union of surviving cells and report $C_t = [\inf S_t, \sup S_t]$.
    if C_{t-1} is unbounded **then**
        Update θ_{t+1} by (15), using the $\widehat{b}_t$ computed above.
    end if
    Update the two cumulative sums in (9) using (r_t, Z_t).
end for

Output: A nonasymptotic $(1-\alpha)$-CS $(C_t)_{t \ge 1}$ for q_τ.

We store only surviving cells' endpoints and capitals, the two sums in (9), and the current tracker, query, lattice, and reserve states. Past reports and capital histories need not be revisited.

C Further theoretical results

This appendix establishes finite-interval acquisition and uniformly bounded working memory for the nested-lattice implementation. All theorem proofs are provided in Appendix E. Throughout the proofs, we take $\theta_1 \in \mathbb{R}$, $h_1 > 0$, $b_0 \ge b_{\min} > 0$, $\kappa \in (0,1)$, and $t_0 \in \{0, 1, \ldots\}$, as specified in Sections 3-4. Let

$$\mathcal{G} = \{q_\tau \in S_t \text{ for all } t \ge 1\}, \tag{C1}$$

so that $\Pr(\mathcal{G}) \ge 1 - \alpha$ by Theorem 1.

C.1 Finite-interval acquisition

In the finite CS acquisition stage of the online tracker in Section 3.2, the lattice spacing remains fixed at $\Delta = \Delta_1 = 2^{\lfloor \log_2(\alpha h_1) \rfloor}$. For a proposed center θ, define

$$\begin{aligned} \widetilde{L}(\theta) &= \Delta \left\lfloor \frac{\theta - h_1}{\Delta} \right\rfloor, \quad \widetilde{U}(\theta) = \Delta \left\lceil \frac{\theta + h_1}{\Delta} \right\rceil, \\ g(\theta) &= \rho \left\{ \frac{F(\widetilde{L}(\theta)) + F(\widetilde{U}(\theta))}{2} - \tau \right\}, \\ \mathcal{Z} &= \{\theta \in \mathbb{R} : g(\theta-) \le 0 \le g(\theta+)\}. \end{aligned} \tag{C2}$$

Here, $g(\theta)$ is the mean centered tracker signal, and the left and right limits accommodate the discontinuities introduced by rounding. Since g is nondecreasing with limits $-\rho\tau$ and $\rho(1-\tau)$, $\mathcal{Z}$ is a nonempty compact interval.

Assumption C1. *There exist $u > 0$ and $d > 0$ such that, whenever $\mathrm{dist}(\theta, \mathcal{Z}) < u$,*

$$\rho \min\left[\tau - F\{\widetilde{L}(\theta)\}, F\{\widetilde{U}(\theta)\} - \tau\right] \geq d. \tag{C3}$$

Remark C1. Near $\mathcal{Z}$, the condition keeps $F(\widetilde{L}(\theta))$ below τ and $F(\widetilde{U}(\theta))$ above τ, each by at least d/ρ. These gaps, together with the prepaid reserve, ensure finite-time rejection of both tails. The condition concerns the chosen h_1 and Δ and requires no density continuity. The algorithm need not know u or d.

Theorem C1. *Under Assumption C1, the procedure obtains a bounded, nonempty CS or terminates with an empty set in finite time almost surely. In particular, the acquisition time in (14) satisfies*

$$\Pr\big(\mathcal{G} \cap \{\tilde{\sigma} = \infty\}\big) = 0, \quad \Pr(\tilde{\sigma} < \infty) \geq 1 - \alpha.$$

Remark C2. On $\mathcal{G}$ in (C1), the true candidate remains in every surviving set, so empty termination is impossible and Stage 2 of the online tracker in Section 3.2 begins in finite time. The initial position may be arbitrary, and the initial query need not contain q_τ. The theorem establishes finite acquisition time without requiring a prespecified bounded parameter space; it does not assert a deterministic acquisition deadline or a finite expected acquisition time.

C.2 Uniformly bounded working memory

Let M_t be the number of surviving cell records after round t, with $M_0 = 1$. We measure memory by the number of stored numerical values.

Theorem C2. *For any fixed $\alpha, \lambda \in (0, 1)$ and any distribution F on $\mathbb{R}$, the nested-lattice implementation satisfies*

$$\sup_{t \geq 0} M_t < \infty, \quad \textit{almost surely}, \quad \mathbb{E}\left(\sup_{t \geq 0} M_t\right) < \infty.$$

The implementation uses $\mathcal{O}(1 + \sup_{t \geq 0} M_t)$ peak working memory.

Remark C3. The working memory remains bounded as the number of users increases, without requiring the acquisition condition or density regularity. In contrast, the standard exact order-statistic implementation of Howard & Ramdas (2022) uses $\mathcal{O}(T)$ memory for T observations. Our method thus provides a more memory-efficient approach to online quantile inference.

D Additional numerical results

Table D1: Median first finite-interval time $\tilde{\sigma}$ and SCP (%) under $\mathcal{N}(0, 1)$. Here $\varepsilon = \infty$ denotes the non-private case; k and m denote 10^3 and 10^6 users.

	$\tau = 0.3$					$\tau = 0.5$					$\tau = 0.7$				
ε	0.5	1	2	3	∞	0.5	1	2	3	∞	0.5	1	2	3	∞
$\tilde{\sigma}$	1,179.5	305.5	129	103	90	899	237	110	91	84	1,217.5	312	127.5	101	91
	SCP (%) at user count t														
50k	99.9	100.0	100.0	99.9	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.9	100.0	99.9	99.9
100k	99.9	100.0	100.0	99.9	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.9	100.0	99.9	99.9
300k	99.9	100.0	100.0	99.9	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.9	100.0	99.9	99.9
1m	99.9	100.0	100.0	99.9	99.9	100.0	100.0	100.0	100.0	100.0	100.0	99.9	100.0	99.9	99.7
5m	99.9	100.0	99.9	99.9	99.9	99.9	100.0	100.0	100.0	100.0	100.0	99.9	100.0	99.9	99.7

D.1 Results for Experiment 1 under alternative distributions

We compare our method with Cai GM and Cai stitched under normal, standard Cauchy, lognormal, and three-component Beta mixture distributions, with $\tau \in \{0.3, 0.5, 0.7\}$, $\varepsilon = 1$, and initial position 0. The normal experiments also assess sensitivity to data scale and the effect of rescaling by a public, prespecified factor. The two asymptotic CSs of Cai et al. (2025) share their estimates and privatized reports, differing only in the confidence boundary.

Normal, Cauchy, and lognormal distributions. Figures D1-D3 present decreasing AW for all three methods, with empirical SCP remaining above the nominal level of 95%. More specifically, our method yields finite intervals substantially earlier, while Cai GM eventually produces narrower intervals. The asymptotic Cai stitched construction and our nonasymptotic method exhibit similar contraction behavior, consistent with their common theoretical rate $\mathcal{O}(\sqrt{\log\log t/t})$. Although our intervals are somewhat wider than those of Cai stitched, our method achieves comparable widths while maintaining nonasymptotic validity. Therefore, our method gives informative quantile inference substantially earlier under the LDP constraint. Additionally, the Cauchy and lognormal results further support the robustness of our method to heavy-tailed and skewed distributions, respectively.

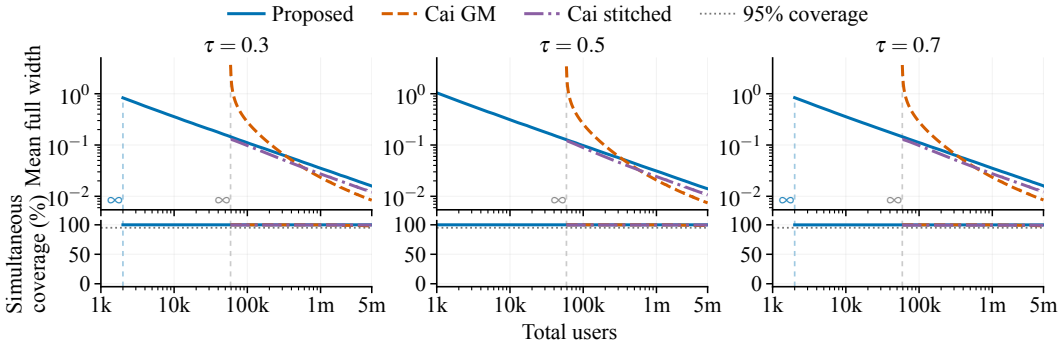

Figure D1: Standard normal distribution $\mathcal{N}(0,1)$: AW (top) and SCP (bottom) at $\varepsilon = 1$. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

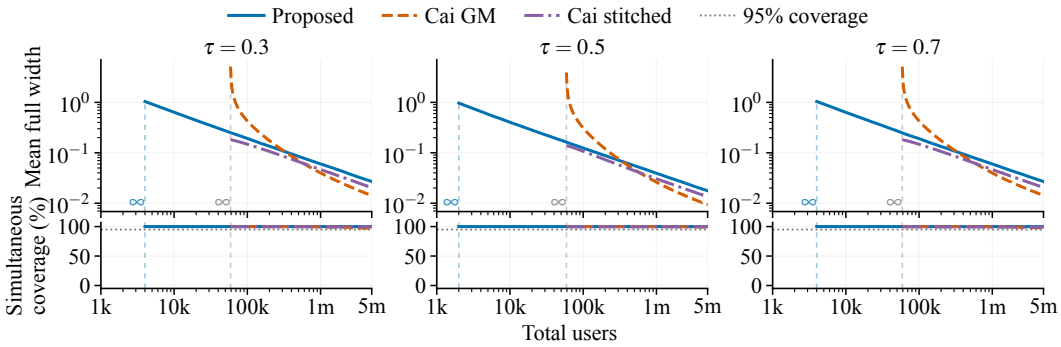

Figure D2: Standard Cauchy distribution $\mathcal{C}(0,1)$: AW (top) and SCP (bottom) at $\varepsilon = 1$. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

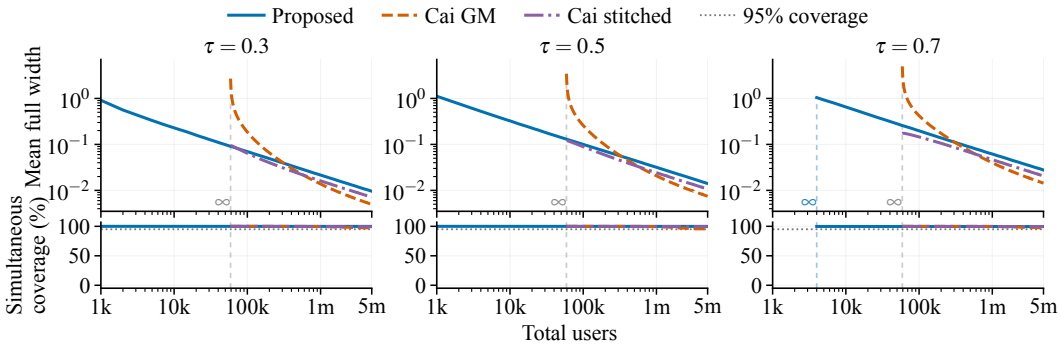

Figure D3: Lognormal distribution $\mathrm{Logn}(0,1)$: AW (top) and SCP (bottom) at $\varepsilon = 1$. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

Three-component Beta mixture distribution. We consider the equally weighted mixture

$$F = \frac{1}{3}\text{Beta}(10, 100) + \frac{1}{3}\text{Beta}(100, 100) + \frac{1}{3}\text{Beta}(100, 10).$$

Figure D4 shows that our method maintains empirical SCP above 95% at all three quantile levels while producing contracting intervals. In contrast, both Cai GM and Cai stitched exhibit severe undercoverage at $\tau = 0.3$ and 0.7, despite decreasing AW. The mixture has a smooth positive density at these target quantiles, so this behavior illustrates finite-sample limitations of the asymptotic constructions under the specified initialization and tuning. These results further demonstrate the value of our nonasymptotic coverage guarantee for multimodal distributions.

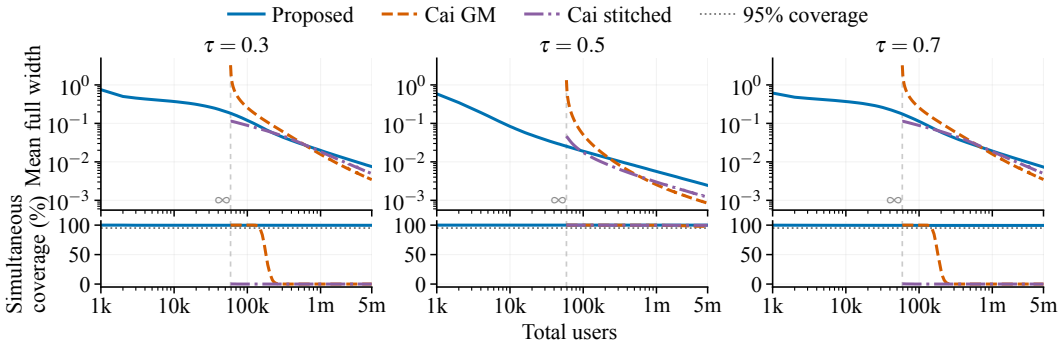

Figure D4: Three-component Beta mixture distribution: AW (top) and SCP (bottom) at $\varepsilon = 1$. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

Sensitivity to data scale. We consider $\mathcal{N}(0, \sigma^2)$ with $\sigma \in \{0.1, 1, 5, 10\}$ and $\tau \in \{0.3, 0.5, 0.7\}$, fixing $\varepsilon = 1$, initial position 0, and the tuning parameters in the main text. Each setting uses 1,000 replications of $T = 5 \times 10^6$ users. Widths are reported in the original data units.

Figure D5 shows that our empirical SCP remains between 99.7% and 100%. Larger scales delay finite-interval acquisition. In contrast, both Cai methods exhibit substantial undercoverage at $\tau = 0.3$ and 0.7: final SCP ranges from 18.9% to 53.2% at $\sigma = 0.1$ and is 0% at $\sigma = 5$ and 10. At $\tau = 0.5$, where the initial position equals the true quantile, all methods maintain empirical SCP above 95%.

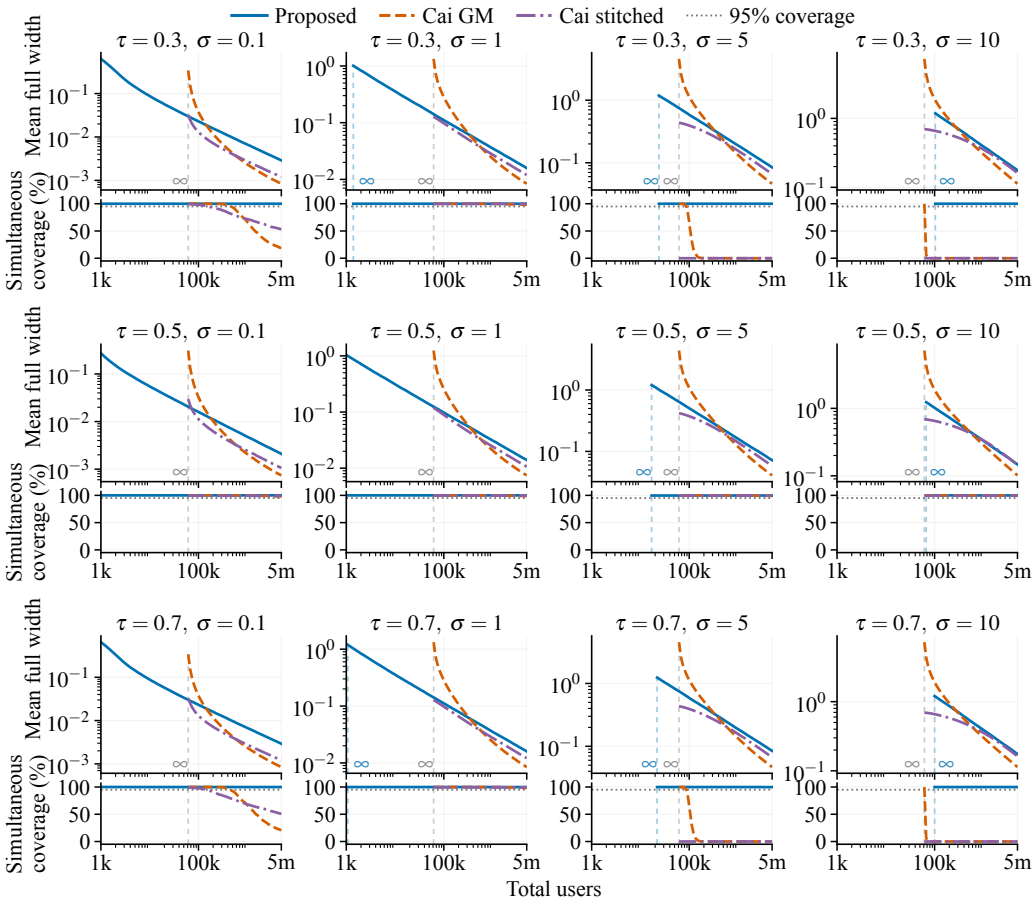

Figure D5: Scale sensitivity with fixed default tuning at $\varepsilon = 1$. Each pair of rows shows AW in the original data units (upper) and SCP (lower). Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

We next apply each method to $Y_t = X_t/s$ for public, prespecified scales with $s/\sigma \in \{0.1, 1, 2\}$, keeping the same numerical tuning and multiplying the reported endpoints by s. For each fixed ratio s/σ, the scaled observations follow the same distribution across all four original scales. With identical tuning, one set of curves therefore represents all four settings. Figure D6 displays these results on the Y_t scale; multiplying widths by s gives the widths in the original units.

For our method, scaling corresponds to a proposed initial radius sh_1 and slope parameters b_0/s and $b_{\min}/s$ in the original units, affecting query localization and betting efficiency without changing coverage validity. For the Cai methods, it corresponds to SGD steps $s\eta_j$ in the original units, with unchanged burn-in; unsuitable scaling can leave location bias large relative to the reported interval widths.

With $s/\sigma = 1$ or 2, our median acquisition times range from 140 to 312, and the final widths in the original units differ by at most 2.7% between these two ratios at each original scale. Both Cai methods attain empirical SCP above 95%. With $s/\sigma = 0.1$, however, our median acquisition times increase to 13,787-19,229, and both Cai methods have 0% SCP at $\tau = 0.3, 0.7$. These results highlight the importance of appropriate standardization for the practical performance of both our method and the Cai methods. Developing scale-invariant tuning rules is a direction for future research.

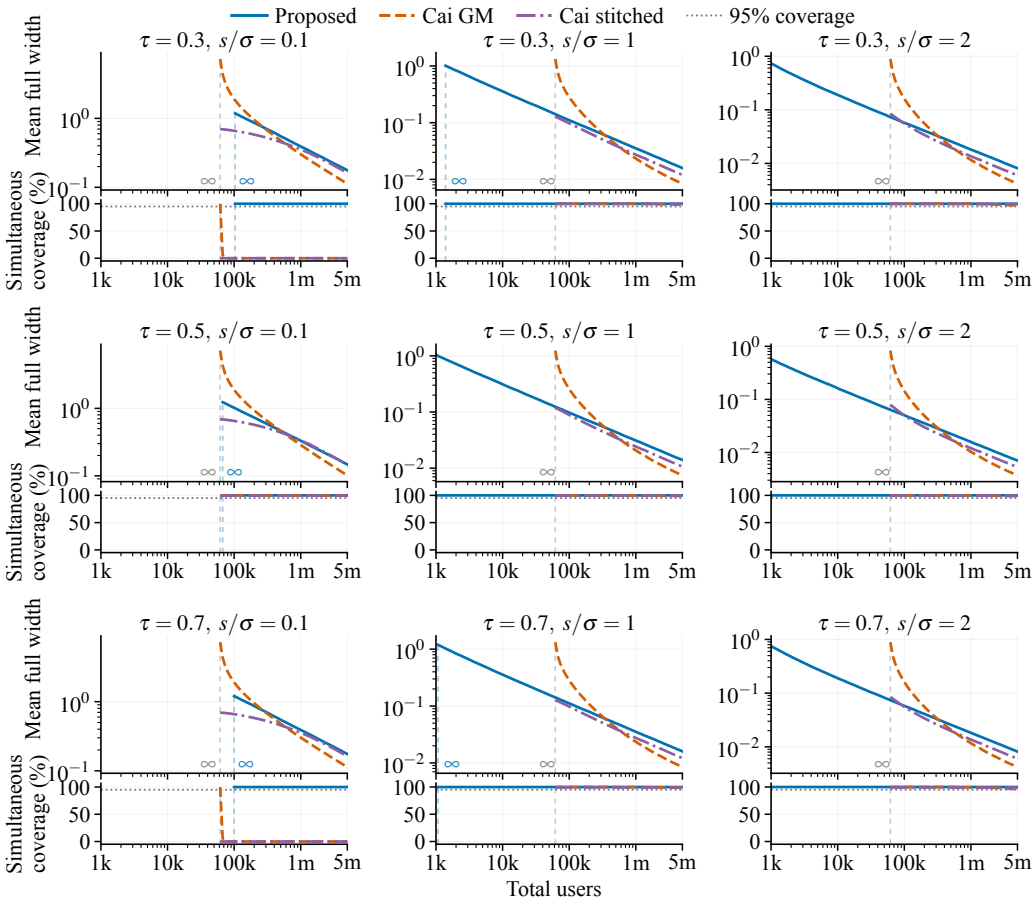

Figure D6: Scale calibration at $\varepsilon = 1$: AW on the $Y_t = X_t/s$ scale (upper) and SCP (lower) for each τ and s/σ. Each column applies to all four original scales; multiply widths by s to recover the original units. Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

D.2 Experiment A1: prepaid capital reserves

We compare the reserve recursion (12) with the direct product $\prod_{s=1}^{t} e_s(q)$ under $\mathcal{N}(0,1)$ at $\tau = 0.5$ and $\varepsilon = 1$. Each setting uses 1,000 replications. Both variants start with unit capital, use the same rejection threshold, and use the tuning parameters in the main text.

The upper row of Figure D7 compares the full adaptive procedures from initial positions 0, 32, and 64, with $T = 5 \times 10^6$. Each curve shows the fraction of replications that obtain a nonempty nonasymptotic CS with full width at most 0.2. The lower row isolates the capital update by testing the false candidate $q^\dagger = 0.5$ with identical queries and reports. We use the query $(0.75, 1.75]$ for the first H users and the query $(-0.25, 0.25]$ thereafter, where $H \in \{0, 1{,}000, 10{,}000\}$, and record rejection through user $H + 50{,}000$ without resetting the capital or slope estimate. Higher curves indicate that more replications have reached the width threshold (top) or rejected the false candidate (bottom); the shaded gaps quantify these differences.

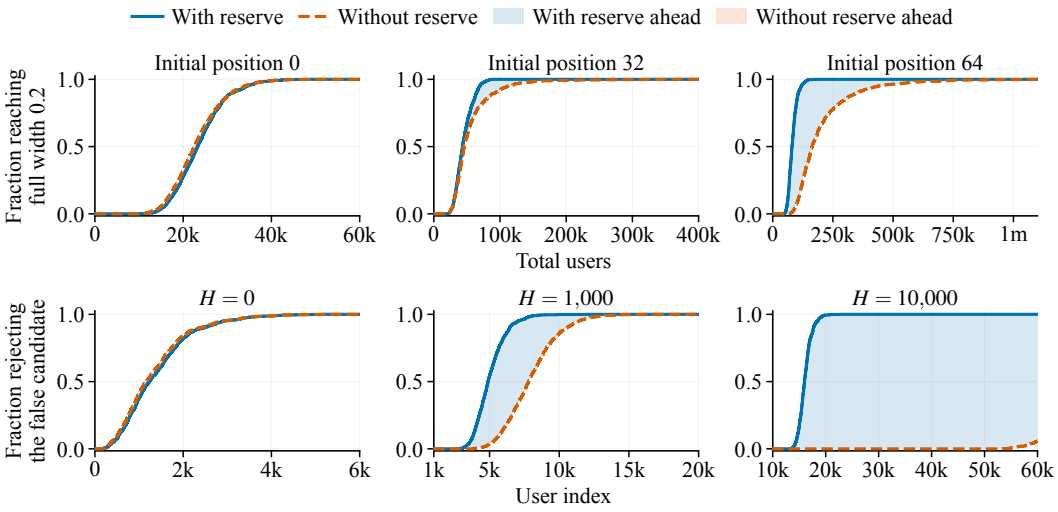

Figure D7: Reserve ablation under $\mathcal{N}(0,1)$ at $\tau = 0.5$ and $\varepsilon = 1$: fractions reaching full width 0.2 (top) and rejecting $q^\dagger = 0.5$ (bottom). Shading shows differences between the curves. Here k and m denote 10^3 and 10^6 users; both axes use linear scales.

For initial position 0 in the upper row and $H = 0$ in the lower row, the two curves are close, and the direct product is slightly ahead over part of the trajectory. The reserve advantage becomes pronounced under distant initialization or prolonged unfavorable queries. At $H = 10{,}000$, all reserve paths reject by user $60{,}000$, compared with 5.8% without reserve. This supports the role of prepaid capital in reducing recovery delays after early betting losses.

D.3 Experiment A2: lattice rounding

We compare implementations with and without lattice rounding under $\mathcal{N}(0,1)$, with $\tau \in \{0.3, 0.5, 0.7\}$, $\varepsilon = 1$, and initial position 0. Each setting uses 100 replications of $T = 5 \times 10^6$ users. We use the reserve and tuning parameters in the main text. For each quantile, Figure D8 shows CS paths with and without lattice rounding from one prespecified replication, with early and late periods displayed separately.

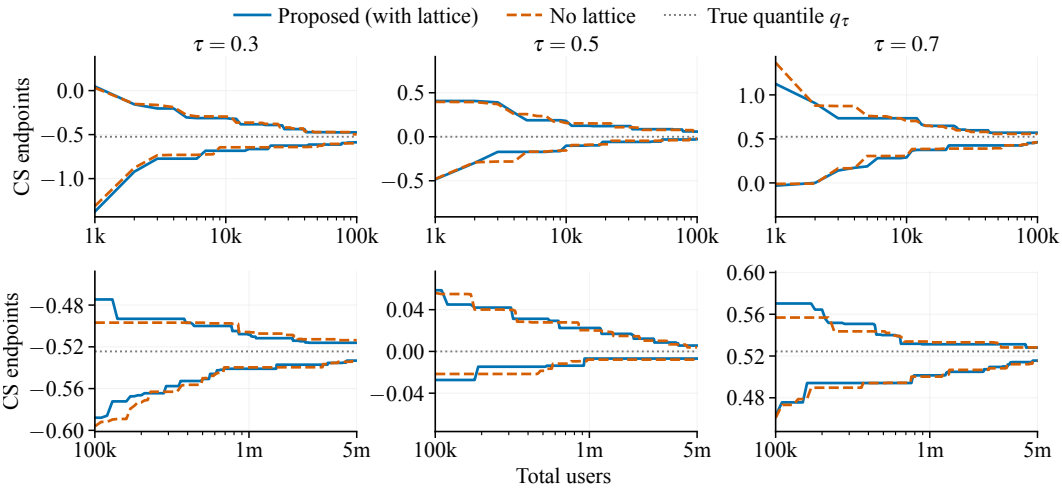

Figure D8: CS trajectories under $\mathcal{N}(0,1)$ with and without lattice rounding at $\varepsilon = 1$. Top: $10^3 \le t \le 10^5$; bottom: $10^5 \le t \le 5 \times 10^6$. Dotted lines mark true quantiles. Here k and m denote 10^3 and 10^6 users; user count uses a log scale.

The paired trajectories contract similarly toward the true quantiles. We also find that across all replications, both methods attain 100% empirical SCP, with comparable final AW. Rounding reduces the final mean cell count from approximately 50-65 to 14, indicating substantial storage savings with little effect on interval width.

D.4 Experiment A3: surviving cells

We examine the number of surviving cells under $\mathcal{N}(0,1)$ at $\tau = 0.5$, using initial position 0 and $\varepsilon \in \{0.5, 1, 2, 3\}$. Each setting uses 1,000 replications of $T = 5 \times 10^6$ users. Figure D9 reports the mean cell count, which determines the storage and update costs of the implementation.

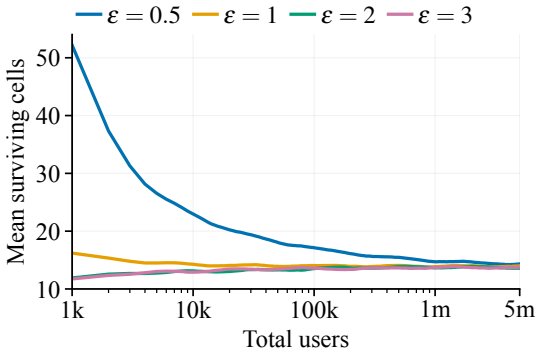

Figure D9: Mean number of surviving cells under $\mathcal{N}(0,1)$ at $\tau = 0.5$ across privacy budgets. Here k and m denote 10^3 and 10^6 users; user count uses a log scale.

Stronger privacy increases the early cell count, particularly at $\varepsilon = 0.5$. As the intervals contract, the counts become similar across privacy budgets and remain near 14 at late times. This behavior supports the storage advantage established in Theorem C2.

D.5 Experiment A4: width scaling

To assess the width scaling in Theorem 2, we use standard normal and Cauchy data with $\tau \in \{0.3, 0.5, 0.7\}$, $\varepsilon \in \{0.5, 1, 2, 3, \infty\}$, and initial position 0. Each setting uses 1,000 replications of $T = 5 \times 10^6$ users. Figures D10-D11 display the normalized AW

$$\rho f(q_\tau)\,\mathrm{AW}(t)\sqrt{\frac{t}{\log\log t}},$$

where $\rho = 1$ in the non-private case. The known density $f(q_\tau)$ is used only for normalization. Vertical dashed lines mark the first plotted checkpoint with finite AW.

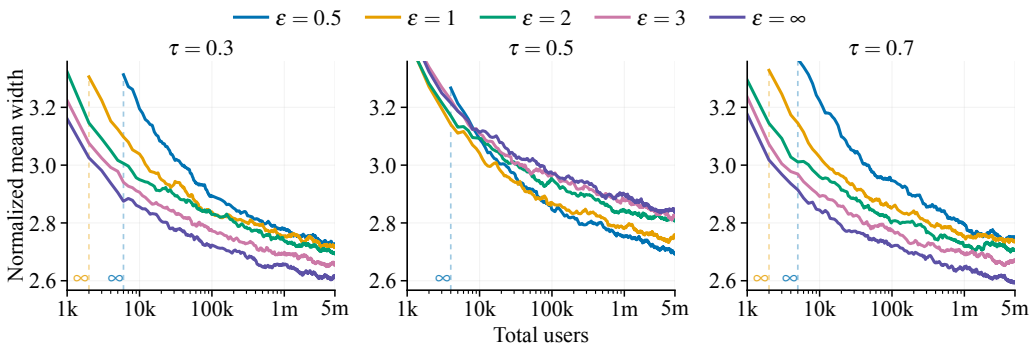

Figure D10: Normalized AW under $\mathcal{N}(0,1)$ across privacy budgets. Here $\varepsilon = \infty$ denotes the non-private case; k and m denote 10^3 and 10^6 users; user count uses a log scale.

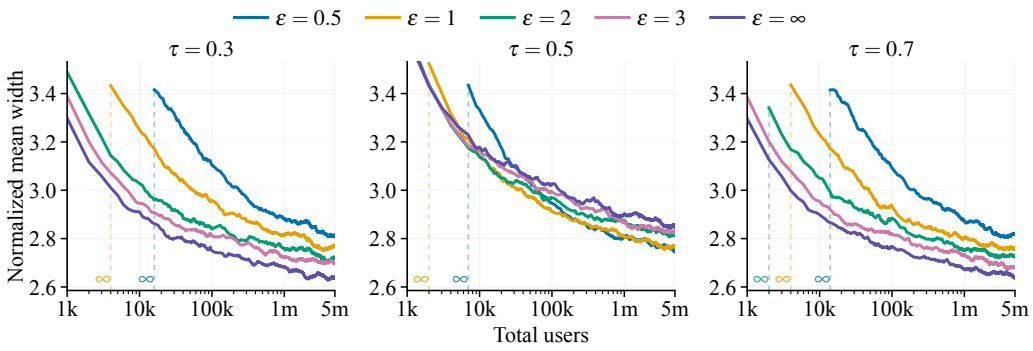

Figure D11: Normalized AW under $\mathcal{C}(0,1)$ across privacy budgets. Here $\varepsilon = \infty$ denotes the non-private case; k and m denote 10^3 and 10^6 users; user count uses a log scale.

The normalized widths decrease gradually and remain on a comparable scale across quantiles and privacy budgets, ranging from approximately 2.6 to 2.9 at T. This pattern is empirically consistent with the theoretical contraction rate suggested by Theorem 2. The Cauchy results also illustrate that heavy tails do not prevent this behavior when the density is positive and smooth near q_τ.

D.6 Experiment A5: tuning parameter sensitivity

We assess sensitivity under $\mathcal{N}(0,1)$ at $\tau \in \{0.3, 0.5, 0.7\}$ and $\varepsilon = 1$, with initial position 0. We vary one parameter at a time over $t_0 \in \{32, 64, 128\}$, $b_0 \in \{0.09, 0.18, 0.36\}$, $b_{\min} \in \{0.005, 0.01, 0.02\}$, $\kappa \in \{0.5, 0.75, 0.9\}$, and $\lambda \in \{0.5, e^{-1/2}, 0.75\}$, holding the others at their default values. Each setting uses 1,000 replications of $T = 5 \times 10^6$ users. When t_0 or b_0 changes, h_1 is recalculated by (B2).

Figures D12-D14 show similar contraction patterns, with final SCP between 99.6% and 100% across all settings. In the settings considered, larger t_0 or b_0 moderately delays finite-interval acquisition but has little effect on late-time AW, while varying $b_{\min}$ produces nearly identical widths. Smaller κ gives slightly wider intervals, and the default $\lambda = e^{-1/2}$ yields the smallest final AW among the three query fractions. These results indicate stable coverage and modest width sensitivity over the tested ranges.

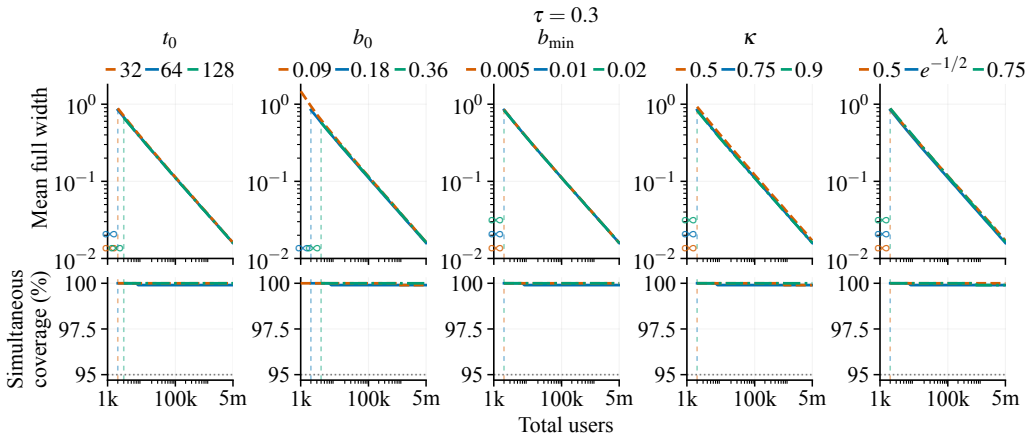

Figure D12: Tuning parameter sensitivity under $\mathcal{N}(0,1)$ at $\tau = 0.3$ and $\varepsilon = 1$: AW (top) and SCP (bottom). Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

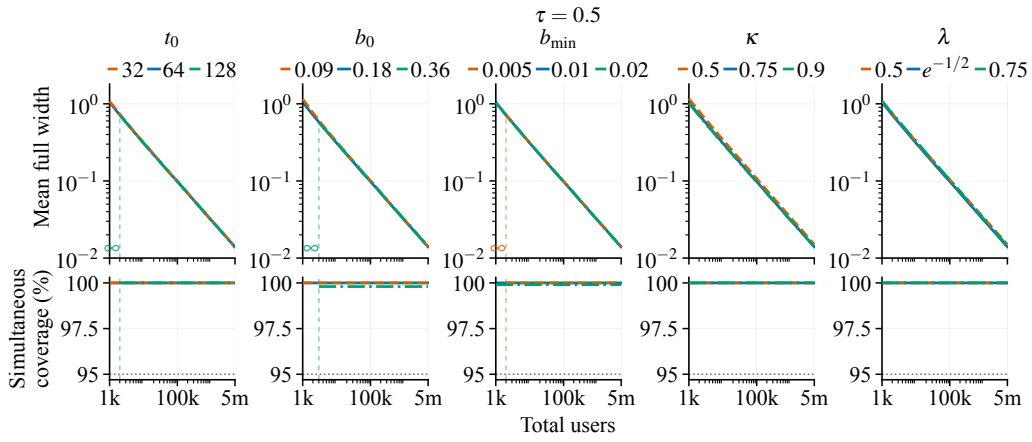

Figure D13: Tuning parameter sensitivity under $\mathcal{N}(0,1)$ at $\tau = 0.5$ and $\varepsilon = 1$: AW (top) and SCP (bottom). Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

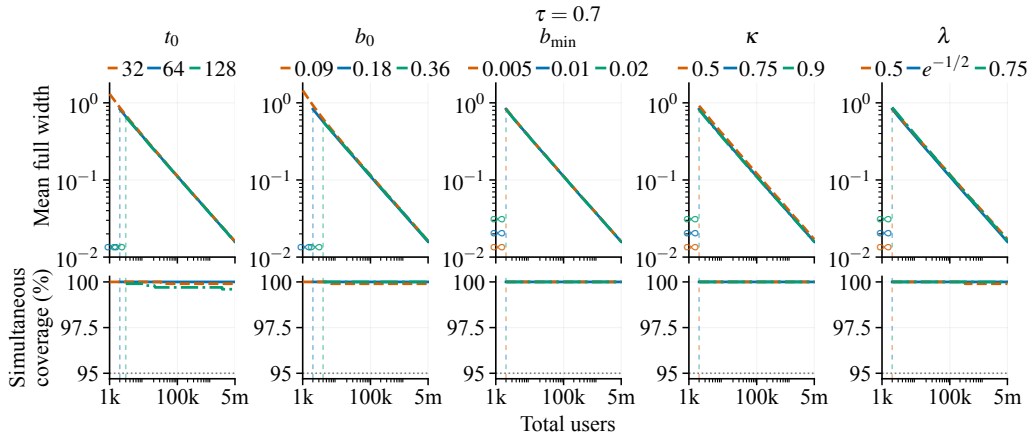

Figure D14: Tuning parameter sensitivity under $\mathcal{N}(0,1)$ at $\tau = 0.7$ and $\varepsilon = 1$: AW (top) and SCP (bottom). Here k and m denote 10^3 and 10^6 users; user count and width use log scales.

E Proofs of the theoretical results

At each round t, the query endpoints and forecasts are $\mathcal{F}_{t-1}$-measurable, and user t releases one report using independent randomization. The proofs analyze the lattice implementation with rounded query endpoints $\widetilde{L}_t, \widetilde{U}_t$. Recall the event $\mathcal{G}$ in (C1), and write $\Lambda(B; v, c) = (v/c)^B\{(1-v)/(1-c)\}^{1-B}$ for the betting factor in (7). After termination with an empty set, we set all subsequent betting factors to one, keeping the capital processes constant.

E.1 Proof of Theorem 1

Proof of Theorem 1. (i) Privacy. Appendix B.1 establishes ε-LDP for each report conditional on the past reports. Each user contributes only one report, after which the algorithm no longer accesses the user's original observation. With all other users' inputs fixed, the preceding reports do not depend on this observation, while subsequent reports and released CSs are randomized post-processing of the preceding reports and the user's privatized report. Thus, the full procedure preserves the same privacy guarantee (Dwork & Roth, 2014, Proposition 2.1).

(ii) Coverage. We first relate coverage to survival, then verify the factor and capital inequalities at the true quantile, and finally apply Ville's inequality.

Survival and the reported interval. Recall S_t in (13) and $C_t = [\inf S_t, \sup S_t]$ for $S_t \neq \varnothing$. The set S_t contains the candidates never rejected through time t, whereas C_t is the reported closed interval containing them. Since $S_t \subseteq C_t$, survival implies simultaneous coverage

$$\mathcal{G} = \bigcap_{t\geq 1} \{q_\tau \in S_t\} \subseteq \bigcap_{t\geq 1} \{q_\tau \in C_t\}. \tag{E1}$$

It therefore suffices to prove $\Pr(\mathcal{G}) \geq 1 - \alpha$.

Conditional factor validity at the true quantile. The quantile definition and right continuity of F imply

$$F(q_\tau -) \leq \tau \leq F(q_\tau), \quad x < q_\tau \Rightarrow F(x) < \tau, \quad x \geq q_\tau \Rightarrow F(x) \geq \tau.$$

For $B \sim \text{Bernoulli}(p)$, the constructed factor obeys

$$\mathbb{E}\Lambda(B; v, c) = \frac{pv}{c} + \frac{(1-p)(1-v)}{1-c} = 1 + \frac{(p-c)(v-c)}{c(1-c)}. \tag{E2}$$

Because $0 < c < 1$, its expectation is at most one whenever $(p - c)(v - c) \leq 0$. Given $\mathcal{F}_{t-1}$, the query endpoints and forecasts are fixed. Recall that the private means at the rounded endpoints are

$$p_{t,L} = a + \rho F(\widetilde{L}_t), \quad p_{t,U} = 2a + \rho F(\widetilde{U}_t),$$

whereas $c_L = a + \rho\tau$ and $c_U = 2a + \rho\tau$. For $t > t_0$, we now check each branch of (11); the first t_0 rounds use the neutral factor one.

If $q_\tau \leq \widetilde{L}_t$, then $F(\widetilde{L}_t) \geq \tau$, so $p_{t,L} \geq c_L$. The forecast construction gives $v_{t,L} < c_L$. Hence,

$$\mathbb{E}(e_{t,L} \mid \mathcal{F}_{t-1}) = \mathbb{E}\{\Lambda(B_{t,L}; v_{t,L}, c_L) \mid \mathcal{F}_{t-1}\} = 1 + \frac{(p_{t,L} - c_L)(v_{t,L} - c_L)}{c_L(1 - c_L)} \leq 1.$$

If $q_\tau > \widetilde{U}_t$, then $F(\widetilde{U}_t) < \tau$, so $p_{t,U} < c_U$. Since $v_{t,U} > c_U$, similarly $\mathbb{E}(e_{t,U} \mid \mathcal{F}_{t-1}) \leq 1$. If $\widetilde{L}_t < q_\tau \leq \widetilde{U}_t$, the selected factor is $e_t(q_\tau) = 1$.

Supermartingale property of $E_t(q_\tau)$. Recall that $R_t = \{1+\log(t+1)\}^{-1}$ is nonnegative and nonincreasing. Starting from $E_0(q_\tau) = R_0 = 1$, suppose that $E_{t-1}(q_\tau) \geq R_{t-1}$. Then, the amount available for betting is nonnegative

$$E_{t-1}(q_\tau) - R_t \geq R_{t-1} - R_t \geq 0.$$

Since the betting factor $e_t(q_\tau)$ is also nonnegative, the update gives

$$E_t(q_\tau) = R_t + \{E_{t-1}(q_\tau) - R_t\} e_t(q_\tau) \geq R_t \geq 0.$$

Thus, by induction, the capital remains nonnegative at every time.

Using $\mathbb{E}\{e_t(q_\tau) \mid \mathcal{F}_{t-1}\} \leq 1$ in the capital recursion yields

$$\begin{aligned}\mathbb{E}\{E_t(q_\tau) \mid \mathcal{F}_{t-1}\} &= R_t + \{E_{t-1}(q_\tau) - R_t\}\mathbb{E}\{e_t(q_\tau) \mid \mathcal{F}_{t-1}\} \\ &\leq R_t + \{E_{t-1}(q_\tau) - R_t\} = E_{t-1}(q_\tau).\end{aligned}$$

Thus, $\{E_t(q_\tau)\}_{t\geq 0}$ is a nonnegative supermartingale with $E_0(q_\tau) = 1$.

Ville's inequality and simultaneous coverage. For every $n \geq 0$, Ville's inequality gives

$$\Pr\left\{\max_{0\leq t\leq n} E_t(q_\tau) \geq 1/\alpha\right\} \leq \alpha\, \mathbb{E}\{E_0(q_\tau)\} = \alpha.$$

Since $\Pr(\mathcal{G}^c) = \lim_{n\to\infty} \Pr\{\max_{0\leq t\leq n} E_t(q_\tau) \geq 1/\alpha\} \leq \alpha$, we obtain $\Pr(\mathcal{G}) \geq 1 - \alpha$. Combining this with (E1) yields

$$\Pr\{q_\tau \in C_t \text{ for all } t \geq 1\} \geq \Pr(\mathcal{G}) \geq 1 - \alpha.$$

On the simultaneous-coverage event, $q_\tau \in C_T$ for every almost surely finite stopping time $T \geq 1$. □

E.2 Proof of Theorem C1

We first establish that the balance set is nonempty and record the range of the actual acquisition radius. Recall g in (C2). During acquisition, the proposed radius is fixed at h_1 and the lattice spacing is $\Delta = 2^{\lfloor \log_2(\alpha h_1) \rfloor}$. The endpoints $\widetilde{L}(\theta)$ and $\widetilde{U}(\theta)$ are nondecreasing and tend to $-\infty$ or $+\infty$ as θ does. Thus, g is nondecreasing, and, for every proper CDF F,

$$\lim_{\theta \to -\infty} g(\theta) = -\rho\tau < 0, \quad \lim_{\theta \to +\infty} g(\theta) = \rho(1-\tau) > 0.$$

Define $z_- := \inf\{\theta : g(\theta) \geq 0\}$ and $z_+ := \sup\{\theta : g(\theta) \leq 0\}$. These thresholds are finite and satisfy $z_- \leq z_+$: otherwise, any $\theta \in (z_+, z_-)$ would satisfy both $g(\theta) > 0$ and $g(\theta) < 0$. By monotonicity,

$$\mathcal{Z} = [z_-, z_+] = \{z : g(z-) \leq 0 \leq g(z+)\},$$

so $\mathcal{Z}$ is a nonempty compact interval.

For an acquisition query centered at θ_t, the outward rounding in (16) gives

$$\theta_t - h_1 - \Delta < \widetilde{L}_t \leq \theta_t - h_1, \quad \theta_t + h_1 \leq \widetilde{U}_t < \theta_t + h_1 + \Delta.$$

Hence the actual query radius satisfies

$$h_1 \leq r_t < h_1 + \Delta. \tag{E3}$$

For the proof, consider a process that follows Stage 1 in Section 3.2 indefinitely with proposed radius h_1 and lattice spacing Δ. It agrees with the actual process until acquisition or empty termination. Below, $\mathcal{F}_t$ denotes the report filtration of this continued Stage 1 process, and expectations are taken under its original law, without conditioning on $\mathcal{G}$.

Lemma E.1 (Convergence of the acquisition tracker)**.** *With stopping and refinement disabled, let $(\theta_t)_{t \geq 1}$ follow the online tracker update in (15) indefinitely, with fixed radius h_1 and lattice spacing Δ. For every proper CDF F, there exists a $\mathcal{Z}$-valued random variable θ_∞ such that*

$$\lim_{t \to \infty} \theta_t = \theta_\infty, \quad \textit{almost surely}.$$

Proof. Throughout this proof, θ_t denotes the acquisition tracker continued indefinitely without stopping or switching to refinement. Fix a deterministic $z \in \mathcal{Z}$. Since θ_t and $\widehat{b}_t$ are $\mathcal{F}_{t-1}$-measurable, expanding the squared update in (15) gives

$$\mathbb{E}\{(\theta_{t+1} - z)^2 \mid \mathcal{F}_{t-1}\} = (\theta_t - z)^2 - \frac{2(\theta_t - z)g(\theta_t)}{(t+t_0)\widehat{b}_t} + \frac{\mathbb{E}\{(K_t - c_0)^2 \mid \mathcal{F}_{t-1}\}}{(t+t_0)^2 \widehat{b}_t^2}. \tag{E4}$$

The definition of K_t implies $\mathbb{E}(K_t - c_0 \mid \mathcal{F}_{t-1}) = g(\theta_t)$. Let $B_K := \max\{c_0, 1 - c_0\}$. Since $K_t \in [0,1]$, we have $|K_t - c_0| \leq B_K$. Together with $\widehat{b}_t \geq b_{\min} > 0$, this bounds the last term in (E4) and yields

$$\mathbb{E}\{(\theta_{t+1} - z)^2 \mid \mathcal{F}_{t-1}\} \leq (\theta_t - z)^2 - \frac{2(\theta_t - z)g(\theta_t)}{(t+t_0)\widehat{b}_t} + \frac{B_K^2}{b_{\min}^2 (t+t_0)^2}. \tag{E5}$$

By the definitions of z_- and z_+ and the monotonicity of g,

$$g(\theta) < 0 \quad \text{if } \theta < z_-, \quad g(\theta) > 0 \quad \text{if } \theta > z_+.$$

For the fixed $z \in \mathcal{Z}$, monotonicity also gives $g(\theta) \leq g(z-) \leq 0$ when $\theta < z$ and $g(\theta) \geq g(z+) \geq 0$ when $\theta > z$. Consequently,

$$\begin{aligned} (\theta - z)g(\theta) &\geq 0 \quad \text{for every } \theta \in \mathbb{R}, \\ (\theta - z)g(\theta) &> 0 \quad \text{whenever } \theta \notin \mathcal{Z}. \end{aligned} \tag{E6}$$

The second term in (E5) is therefore nonpositive. Thus, the conditional expected squared distance is bounded by the current squared distance and the deterministic error $B_K^2 / \{b_{\min}^2 (t+t_0)^2\}$.

We add the sum of the remaining error bounds to the squared distance and define

$$V_t(z) := (\theta_t - z)^2 + \frac{B_K^2}{b_{\min}^2} \sum_{s=t}^{\infty} (s + t_0)^{-2}.$$

Substituting the definition of $V_t(z)$ into (E5) yields

$$\mathbb{E}\{V_{t+1}(z) \mid \mathcal{F}_{t-1}\} \le V_t(z) - \frac{2(\theta_t - z)g(\theta_t)}{(t + t_0)\widehat{b}_t} \le V_t(z), \tag{E7}$$

where the last inequality follows from (E6). Moreover, $V_t(z)$ is nonnegative and $\mathcal{F}_{t-1}$-measurable. The deterministic initial value and the bounds on $K_t - c_0$ and $\widehat{b}_t$ give a deterministic bound on $|\theta_t|$ at every finite time, ensuring integrability. Therefore, $\{V_t(z)\}_{t\ge1}$ is a nonnegative supermartingale with respect to $(\mathcal{F}_{t-1})_{t\ge1}$. By the nonnegative supermartingale convergence theorem (Durrett, 2019, Theorem 4.2.12) and the vanishing deterministic tail sum, there exists an almost surely finite nonnegative limit $\ell_z \in \mathbb{R}$ such that

$$\lim_{t\to\infty} (\theta_t - z)^2 = \ell_z, \quad \text{almost surely.}$$

If $\ell_z = 0$, $\lim_{t\to\infty} \theta_t = \theta_\infty$ where $\theta_\infty = z$. If $\ell_z > 0$, the update gives $|\theta_{t+1} - \theta_t| \le B_K/\{b_{\min}(t + t_0)\} \to 0$. Together with $|\theta_t - z| \to \sqrt{\ell_z} > 0$, this implies that $\theta_t - z$ eventually has a constant sign. Hence, $\lim_{t\to\infty} \theta_t = \theta_\infty$ where $\theta_\infty \in \{z - \sqrt{\ell_z}, z + \sqrt{\ell_z}\}$. Suppose that $\theta_\infty \notin \mathcal{Z}$. Since $\theta_t \to \theta_\infty$, monotonicity of g and (E6) imply that, for some $c > 0$,

$$(\theta_t - z)g(\theta_t) \ge c$$

for all sufficiently large t. Recall the slope estimator in (9). Since $I(Z_s = 1) - a \le 1 - a$ and $r_s \ge h_1$ by (E3), for $t \ge 2$,

$$\frac{\sum_{s<t} r_s\{I(Z_s = 1) - a\}}{2\sum_{s<t} r_s^2} \le \frac{(1 - a)\sum_{s<t} r_s}{2\sum_{s<t} r_s^2} \le \frac{1 - a}{2h_1}.$$

Including the initialization $\widehat{b}_1 \ge b_{\min}$ gives

$$b_{\min} \le \widehat{b}_t \le B_h := \max\left\{\widehat{b}_1, \frac{1 - a}{2h_1}\right\}, \quad t \ge 1.$$

Consequently,

$$\sum_{t=1}^{\infty} \frac{2(\theta_t - z)g(\theta_t)}{(t + t_0)\widehat{b}_t} \ge \frac{2}{B_h} \sum_{t=1}^{\infty} \frac{(\theta_t - z)g(\theta_t)}{t + t_0} = \infty, \tag{E8}$$

because of $\sum_{t\ge1}(t + t_0)^{-1} = \infty$. Taking expectations in (E7) and summing over $t = 1, \dots, n$ give

$$\mathbb{E}\left\{\sum_{t=1}^{n} \frac{2(\theta_t - z)g(\theta_t)}{(t + t_0)\widehat{b}_t}\right\} \le V_1(z) - \mathbb{E}\{V_{n+1}(z)\} \le V_1(z).$$

By monotone convergence, the infinite sum has expectation at most $V_1(z) < \infty$ and is therefore almost surely finite. This contradicts the result in (E8). Hence, $\theta_\infty \in \mathcal{Z}$, almost surely. □

The following lemma shows how the prepaid reserve supports threshold crossing when the conditional mean betting factors remain uniformly greater than one.

Lemma E.2 (Reserve-funded evidence)**.** *Fix a deterministic integer $s \ge 0$. Let $(e_t)_{t>s}$ be adapted to $(\mathcal{F}_t)_{t\ge s}$ and satisfy $0 \le e_t \le C$, almost surely, for a deterministic $C \ge 2$. Recall the reserve releases $\beta_t = R_{t-1} - R_t$ used in (B1), and define*

$$\underline{A}_s = 0, \quad \underline{A}_t = (\underline{A}_{t-1} + \beta_t)e_t, \quad t > s.$$

Then

$$\limsup_{t\to\infty} \underline{A}_t = \infty, \quad \textit{almost surely on } \left\{\liminf_{t\to\infty} \mathbb{E}(e_t \mid \mathcal{F}_{t-1}) > 1\right\}.$$

Proof. For $0 \le x \le C$ and $0 < \nu < 1$,

$$x^\nu \ge x - (1-\nu)C\log C.$$

Indeed, $x^\nu \ge x$ for $x \le 1$, while for $x > 1$,

$$x^\nu = x\exp\{-(1-\nu)\log x\} \ge x - (1-\nu)x\log x.$$

For each integer $k \ge 1$, choose

$$\nu_k = 1 - \min\left\{\frac{1}{2}, \frac{1}{2kC\log C}\right\} \in [1/2, 1), \quad \delta_k = \frac{1}{2k}.$$

Taking conditional expectations gives

$$m_{t,k} := \mathbb{E}(e_t^{\nu_k} \mid \mathcal{F}_{t-1}) \ge \mathbb{E}(e_t \mid \mathcal{F}_{t-1}) - \frac{1}{2k}. \tag{E9}$$

Choose versions of these conditional expectations so that (E9) holds simultaneously for all t, k outside a single null set. For the reserve in (12),

$$\beta_t = \int_t^{t+1} \frac{dx}{x(1+\log x)^2}, \quad 0 < \beta_t < 1, \quad \beta_t \sim \frac{1}{t(\log t)^2}.$$

Hence $\sum_{t=1}^\infty \beta_t^{\nu_k} = \infty$ for every k.

Fix integers $k, H \ge 1$ and $m > s$. Start an auxiliary account at $Q_{m-1} = 0$ and set

$$I_t = I(Q_{t-1} < H)\prod_{u=m}^{t} I(m_{u,k} \ge 1 + \delta_k),$$
$$Q_t = I_t(Q_{t-1} + \beta_t)e_t + (1 - I_t)Q_{t-1}, \quad t \ge m.$$

The indicator I_t is $\mathcal{F}_{t-1}$-measurable. The account remains fixed after reaching H or failure of a conditional moment bound. Since $0 \le Q_t \le C(H+1)$, all expectations below are finite. Conditional on $\mathcal{F}_{t-1}$,

$$\begin{aligned}\mathbb{E}(Q_t^{\nu_k} - Q_{t-1}^{\nu_k} \mid \mathcal{F}_{t-1}) &= I_t\{(Q_{t-1} + \beta_t)^{\nu_k} m_{t,k} - Q_{t-1}^{\nu_k}\} \\ &\ge I_t\{(1+\delta_k)(Q_{t-1} + \beta_t)^{\nu_k} - Q_{t-1}^{\nu_k}\} \\ &\ge \delta_k I_t \beta_t^{\nu_k}.\end{aligned}$$

The last inequality uses $(Q+\beta)^{\nu_k} \ge Q^{\nu_k}$ and $(Q+\beta)^{\nu_k} \ge \beta^{\nu_k}$ for $Q, \beta \ge 0$. Taking expectations, summing, and applying monotone convergence give

$$\delta_k\,\mathbb{E}\left(\sum_{t=m}^\infty I_t\beta_t^{\nu_k}\right) \le \{C(H+1)\}^{\nu_k} < \infty.$$

Therefore,

$$\sum_{t=m}^\infty I_t\beta_t^{\nu_k} < \infty, \quad \text{almost surely.} \tag{E10}$$

Take the countable intersection of these probability-one events over all integers $k, H \ge 1$ and $m > s$, together with the event on which (E9) holds for all t, k. On a path in this common event with $\liminf_{t\to\infty} \mathbb{E}(e_t \mid \mathcal{F}_{t-1}) > 1$, choose k such that, for all sufficiently large t,

$$\mathbb{E}(e_t \mid \mathcal{F}_{t-1}) \ge 1 + \frac{1}{k}, \quad m_{t,k} \ge 1 + \delta_k.$$

Fix any integer threshold $H \ge 1$ and integer time $n_0 \ge s$. Choose an integer $m > n_0$ after the displayed bounds hold permanently. If Q_t never reaches H, then $I_t = 1$ for all $t \ge m$, contradicting (E10) and $\sum_{t=m}^\infty \beta_t^{\nu_k} = \infty$. Thus, Q_t first reaches H at some finite time $t_* \ge m$. Before and at this crossing, the accounts follow the same recursion, so

$$\underline{A}_t - Q_t = \underline{A}_{m-1}\prod_{u=m}^{t} e_u \ge 0, \quad m \le t \le t_*.$$

In particular, $\underline{A}_{t_*} \ge H$. Since n_0 and H were arbitrary, $\limsup_{t\to\infty} \underline{A}_t = \infty$. The common event was constructed for all deterministic k, m, H before these indices were selected; the argument does not require the time after which the bounds hold to be a stopping time. □

Proof of Theorem C1. We continue the Stage 1 process from Lemma E.1 indefinitely; it agrees with the actual procedure until acquisition or empty termination. Parts (i)-(ii) concern this continued Stage 1 process, and part (iii) returns to the actual procedure.

(i) Conditional mean growth of the betting factors. Lemma E.1 gives $\mathrm{dist}(\theta_t, \mathcal{Z}) \to 0$, almost surely. Hence, (C3) implies that, almost surely, for all sufficiently large t,

$$c_L - p_{t,L} = \rho[\tau - F\{\widetilde{L}(\theta_t)\}] \geq d, \quad p_{t,U} - c_U = \rho[F\{\widetilde{U}(\theta_t)\} - \tau] \geq d.$$

Define the positive constants

$$s_L = c_L \left\{1 - \exp\left(-\frac{\kappa b_{\min} h_1}{c_L}\right)\right\},$$
$$s_U = (1 - c_U) \left\{1 - \exp\left(-\frac{\kappa b_{\min} h_1}{1 - c_U}\right)\right\}.$$

Since $\widehat{b}_t \geq b_{\min}$ and $r_t \geq h_1$, (10) gives

$$c_L - v_{t,L} \geq s_L > 0, \quad v_{t,U} - c_U \geq s_U > 0.$$

Using (E2), we obtain, almost surely, for all sufficiently large $t > t_0$,

$$\mathbb{E}(e_{t,j} \mid \mathcal{F}_{t-1}) \geq 1 + \frac{d s_j}{c_j(1 - c_j)} \geq 1 + \eta, \quad j \in \{L, U\},$$
$$\eta = \min_{j \in \{L,U\}} \frac{d s_j}{c_j(1 - c_j)} > 0.$$

The factor definition also gives the deterministic bound

$$0 < e_{t,j} \leq C := \max_{j \in \{L,U\}} \{c_j^{-1}, (1 - c_j)^{-1}\} < \infty, \quad t \geq 1,$$

including the neutral factors for $t \leq t_0$, with $C \geq 2$.

(ii) Reserve-funded threshold crossings. For each $j \in \{L, U\}$, define an auxiliary capital that always receives $e_{t,j}$:

$$H_{0,j} = 1, \quad H_{t,j} = R_t + (H_{t-1,j} - R_t) e_{t,j}, \quad t \geq 1.$$

Its active capital $A_{t,j} := H_{t,j} - R_t$ satisfies

$$A_{0,j} = 0, \quad A_{t,j} = (A_{t-1,j} + \beta_t) e_{t,j} \geq 0.$$

Part (i) verifies the conditions of Lemma E.2. Applying it with $s = 0$ and $e_t = e_{t,j}$ gives

$$\limsup_{t \to \infty} A_{t,j} = \limsup_{t \to \infty} H_{t,j} = \infty, \quad j \in \{L, U\}, \quad \text{almost surely},$$

where the equality follows from $R_t \to 0$.

(iii) Tail elimination and finite acquisition. For $j \in \{L, U\}$, define

$$T_j := \inf\{t \geq 1 : H_{t,j} \geq \alpha^{-1}\}.$$

Since $H_{0,j} = 1 < \alpha^{-1}$ and $\limsup_{t \to \infty} H_{t,j} = \infty$, almost surely, both T_L and T_U are almost surely finite.

We first consider the left tail. Fix any candidate

$$q \leq \min_{1 \leq s \leq T_L} \widetilde{L}_s.$$

Then, $q \leq \widetilde{L}_s$ for every $1 \leq s \leq T_L$, so (11) gives

$$e_s(q) = e_{s,L}, \quad 1 \leq s \leq T_L.$$

Moreover, $E_0(q) = H_{0,L} = 1$. Since $E_s(q)$ and $H_{s,L}$ receive the same factors and satisfy the same reserve recursion, induction gives

$$E_s(q) = H_{s,L}, \quad 0 \leq s \leq T_L.$$

In particular,

$$E_{T_L}(q) = H_{T_L,L} \geq \alpha^{-1}.$$

Therefore, $q \notin S_{T_L}$ by (13). Thus, every candidate in

$$\left(-\infty, \min_{1\leq s\leq T_L} \widetilde{L}_s\right]$$

has been eliminated by time T_L.

Similarly, every candidate $q > \max_{1\leq s\leq T_U} \widetilde{U}_s$ receives $e_{s,U}$ for every $1 \leq s \leq T_U$. Hence,

$$E_{T_U}(q) = H_{T_U,U} \geq \alpha^{-1},$$

so $q \notin S_{T_U}$.

Let $T := T_L \vee T_U$. Since candidate elimination is permanent,

$$S_T \subseteq \left(\min_{1\leq s\leq T_L} \widetilde{L}_s, \max_{1\leq s\leq T_U} \widetilde{U}_s\right].$$

Both endpoints on the right are finite almost surely because $T_L, T_U < \infty$ and every finite-time query has finite endpoints. Therefore, at time T, the continued Stage 1 process has either a bounded nonempty surviving hull or an empty surviving set.

The actual procedure either leaves acquisition before T or agrees with the continued Stage 1 process through time T. It therefore acquires a bounded nonempty interval or terminates with the empty set in finite time almost surely.

Finally, on $\mathcal{G}$, the true candidate satisfies $q_\tau \in S_t$ for every $t \geq 1$, so empty termination is impossible. Consequently,

$$\Pr[\mathcal{G} \cap \{\tilde{\sigma} = \infty\}] = 0.$$

Together with $\Pr(\mathcal{G}) \geq 1 - \alpha$ from Theorem 1, this gives

$$\Pr(\tilde{\sigma} < \infty) \geq \Pr(\mathcal{G}) \geq 1 - \alpha.$$

□

A sufficient constant-density example. The endpoint-gap condition (C3) can be verified as follows. Suppose $F(q_\tau) = \tau$ and $f = f_0 > 0$ on $[q_\tau - h_1 - 2\Delta, q_\tau + h_1 + 2\Delta]$. The lattice rule gives $\Delta \leq \alpha h_1 < h_1$. Let $m(\theta) = \{\widetilde{L}(\theta) + \widetilde{U}(\theta)\}/2$ be the rounded query midpoint. Outward rounding gives $|m(\theta) - \theta| < \Delta/2$. For $|\theta - q_\tau| \leq \Delta$, both endpoints lie in the constant-density interval, so

$$g(\theta) = \rho f_0\{m(\theta) - q_\tau\}.$$

Thus $g(\theta) < 0$ for $q_\tau - \Delta \leq \theta < q_\tau - \Delta/2$, and $g(\theta) > 0$ for $q_\tau + \Delta/2 < \theta \leq q_\tau + \Delta$. By monotonicity and the definition of $\mathcal{Z}$,

$$\mathcal{Z} \subseteq [q_\tau - \Delta/2, q_\tau + \Delta/2].$$

Set $u = \Delta/2$. If $\operatorname{dist}(\theta, \mathcal{Z}) < u$, then $|\theta - q_\tau| < \Delta$, so both endpoints remain in the constant-density interval and

$$\widetilde{L}(\theta) \leq \theta - h_1 < q_\tau - (h_1 - \Delta),$$
$$\widetilde{U}(\theta) \geq \theta + h_1 > q_\tau + (h_1 - \Delta).$$

Consequently,

$$\rho \min[\tau - F\{\widetilde{L}(\theta)\}, F\{\widetilde{U}(\theta)\} - \tau] \geq \rho f_0(h_1 - \Delta) > 0.$$

Taking $d = \rho f_0(h_1 - \Delta)$ verifies Assumption C1.

E.3 Proof of Theorem 2

Return to the actual procedure and its report filtration $\mathcal{F}_t$. Recall $W_t = C_{U,t} - C_{L,t}$ and the acquisition time $\tilde{\sigma}$ in (14). We set $W_t = 0$ after empty termination. For $x > 0$, write $\psi(x) = 2^{\lfloor \log_2 x \rfloor}$ for the dyadic grid rule.

Lattice geometry. At every finite time, a nonempty surviving set contains a cell of positive length; hence its hull has positive width. Recall that, for a bounded, nonempty C_{t-1} with $W_{t-1} > 0$, the Stage 2 query updates are

$$\mu_t = \frac{C_{L,t-1} + C_{U,t-1}}{2}, \quad h_t = \frac{\lambda W_{t-1}}{2}.$$

Outward rounding and $\Delta_t \le \alpha h_t$ give

$$\mu_t - (1+\alpha)h_t \le \mu_t - h_t - \Delta_t < \widetilde{L}_t \le \mu_t - h_t,$$
$$\mu_t + h_t \le \widetilde{U}_t < \mu_t + h_t + \Delta_t \le \mu_t + (1+\alpha)h_t.$$

Since $\lambda < 1/(1+\alpha)$ implies $(1+\alpha)\lambda < 1$,

$$C_{L,t-1} < \widetilde{L}_t < \widetilde{U}_t < C_{U,t-1}, \quad \frac{\lambda W_{t-1}}{2} \le r_t < \frac{(1+\alpha)\lambda W_{t-1}}{2}. \tag{E11}$$

The monotonicity of W_t and $x/2 < \psi(x) \le x$ yield

$$\Delta_t = \min\{\Delta, \psi(\alpha\lambda W_{t-1}/2)\} \ge \min\{\Delta, \alpha\lambda W_{t-1}/4\}. \tag{E12}$$

All historical endpoints belong to the current dyadic lattice. Hence

$$C_t \ne \varnothing,\ W_t < W_{t-1} \implies W_{t-1} - W_t \ge \Delta_t,$$
$$C_t = C_{t-1} \implies (\widetilde{L}_{t+1}, \widetilde{U}_{t+1}, \Delta_{t+1}) = (\widetilde{L}_t, \widetilde{U}_t, \Delta_t).$$

Reserve capital. Recall $\beta_t = R_{t-1} - R_t$, and let $A_t(q) := E_t(q) - R_t$ denote the active capital. The recursion in (12) gives $A_0(q) = 0$ and

$$A_t(q) = \{A_{t-1}(q) + \beta_t\} e_t(q) \ge 0,$$
$$\beta_t = \int_t^{t+1} \frac{dx}{x(1+\log x)^2}, \quad \lim_{t\to\infty} t(\log t)^2 \beta_t = 1.$$

In particular, $0 < \beta_t < 1$ for $t \ge 1$. For analysis, deleted candidates may continue this recursion along the same query path.

Frozen-query consequence. Fix a deterministic report time $s \ge 0$ at which the next query is defined, and fix $j \in \{L, U\}$. Consider a virtual continuation that uses $(\widetilde{L}_{s+1}, \widetilde{U}_{s+1}]$ at every report $t > s$, with the reporting mechanism in (4) and the side factors in (7). All quantities in this paragraph refer to this continuation. Suppose $v_{t,j} \to v_{\infty,j} \in (0,1)$, almost surely, where $v_{\infty,j}$ is $\mathcal{F}_s$-measurable. Apply Lemma E.2 to

$$\underline{A}_s = 0, \quad \underline{A}_t = (\underline{A}_{t-1} + \beta_t) e_{t,j}, \quad t > s.$$

Because the query is fixed and each report uses a fresh user, $\Pr(B_{t,j} = 1 \mid \mathcal{F}_{t-1}) = p_{s+1,j}$. By (E2),

$$\lim_{t\to\infty} \mathbb{E}(e_{t,j} \mid \mathcal{F}_{t-1}) = 1 + \frac{(p_{s+1,j} - c_j)(v_{\infty,j} - c_j)}{c_j(1-c_j)}, \quad \text{almost surely.}$$

Moreover, $0 < e_{t,j} \le \max\{c_j^{-1}, (1-c_j)^{-1}\}$. The lemma therefore gives $\limsup_{t\to\infty} \underline{A}_t = \infty$, almost surely on $\{(p_{s+1,j} - c_j)(v_{\infty,j} - c_j) > 0\}$. Since this event is $\mathcal{F}_s$-measurable,

$$\Pr\left(\limsup_{t\to\infty} \underline{A}_t = \infty \,\middle|\, \mathcal{F}_s\right) = 1, \quad \text{almost surely on } \{(p_{s+1,j} - c_j)(v_{\infty,j} - c_j) > 0\}.$$

The common account provides a lower bound that is uniform over all candidates on the corresponding side of the frozen query. Indeed, for every $q \in S_s$ with $q \le \widetilde{L}_{s+1}$ when $j = L$, or $q > \widetilde{U}_{s+1}$ when $j = U$, the candidate account and the common account receive the same factors and reserve releases. Since $e_u(q) = e_{u,j}$ for every $u > s$, the two recursions give

$$\begin{aligned} A_t(q) - \underline{A}_t &= \{A_{t-1}(q) + \beta_t\} e_{t,j} - (\underline{A}_{t-1} + \beta_t) e_{t,j} \\ &= \{A_s(q) - \underline{A}_s\} \prod_{u=s+1}^{t} e_{u,j} \\ &= A_s(q) \prod_{u=s+1}^{t} e_{u,j} \ge 0, \quad t > s, \end{aligned} \tag{E13}$$

where the last equality uses $\underline{A}_s = 0$, and the inequality follows from $A_s(q) \geq 0$ and the nonnegativity of the factors. Thus, $A_t(q) \geq \underline{A}_t$ simultaneously for all such candidates. Under the strict-alternative condition above, $\limsup_{t\to\infty} \underline{A}_t = \infty$, almost surely, so all these candidates are rejected by a common finite time.

Proof of Theorem 2. Write $b_* := \rho f(q_\tau) > 0$ for the local private slope. Constants $C, c_1, c_2 > 0$ depend only on $\alpha, \tau, \varepsilon, \kappa, \lambda$; C may increase between displays.

Step 1: localization. Fix a deterministic s with bounded nonempty C_s and condition on $\mathcal{F}_s$. Freeze $(\widetilde{L}_{s+1}, \widetilde{U}_{s+1}]$ in a virtual continuation. The conditional strong law and (9) yield

$$\lim_{t\to\infty} \frac{1}{t-s} \sum_{u=s+1}^{t} \{I(Z_u = 1) - a\} = \rho\{F(\widetilde{U}_{s+1}) - F(\widetilde{L}_{s+1})\},$$

$$\lim_{t\to\infty} \widehat{b}_t = \max\left[b_{\min}, \frac{\rho\{F(\widetilde{U}_{s+1}) - F(\widetilde{L}_{s+1})\}}{2r_{s+1}}\right] \in (0, \infty), \quad \text{almost surely.}$$

Thus, (10) implies

$$v_{\infty,L} := \lim_{t\to\infty} v_{t,L} \in (0, c_L), \quad v_{\infty,U} := \lim_{t\to\infty} v_{t,U} \in (c_U, 1), \quad \text{almost surely.}$$

Assumption 1 and monotonicity imply that q_τ is the unique solution to $F(x) = \tau$. Since $\widetilde{L}_{s+1} < \widetilde{U}_{s+1}$, it follows that

$$F(\widetilde{L}_{s+1}) < \tau \quad \text{or} \quad F(\widetilde{U}_{s+1}) > \tau.$$

In the left case,

$$p_{s+1,L} - c_L = \rho\{F(\widetilde{L}_{s+1}) - \tau\} < 0.$$

Together with $v_{\infty,L} < c_L$, this gives $(p_{s+1,L} - c_L)(v_{\infty,L} - c_L) > 0$, so the frozen-query consequence of Lemma E.2 yields

$$\limsup_{t\to\infty} \underline{A}_t = \infty, \quad \text{almost surely.}$$

By (E13), $A_t(q) \geq \underline{A}_t$ simultaneously for all surviving candidates on the left side; hence all are rejected by a common finite time. The right case follows identically from $p_{s+1,U} - c_U = \rho\{F(\widetilde{U}_{s+1}) - \tau\} > 0$ and $v_{\infty,U} > c_U$. By (E11), deleting either outer side strictly changes the hull or makes it empty. The actual and frozen processes agree until the first hull change. Hence, for each deterministic s,

$$\Pr(C_t = C_s \text{ for all } t \geq s \mid \mathcal{F}_s) = 0 \quad \text{on } \{C_s \neq \varnothing,\ 0 < W_s < \infty\}.$$

Taking a countable union of these null events over s, almost surely no bounded hull of positive width remains unchanged forever.

On $\mathcal{G}$, the true candidate q_τ remains in S_t at every report, so the surviving set never becomes empty. Theorem C1 therefore implies that acquisition occurs in finite time, giving $\tilde{\sigma} < \infty$, almost surely on $\mathcal{G}$. Since the hulls are nested, W_t is nonincreasing and $W_\infty := \lim_{t\to\infty} W_t$ exists. Suppose $W_\infty > 0$. The mesh bound and $W_{t-1} \geq W_\infty$ give

$$\Delta_t \geq \min\{\Delta, \alpha\lambda W_{t-1}/4\} \geq \min\{\Delta, \alpha\lambda W_\infty/4\} > 0, \quad t > \tilde{\sigma}.$$

Thus the same positive lower bound applies to $\inf_{t>\tilde{\sigma}} \Delta_t$.

All hull endpoints belong to the current lattice. By nesting, their inward movements are nonnegative multiples of Δ_t, so

$$\begin{aligned} W_{t-1} - W_t &= (C_{L,t} - C_{L,t-1}) + (C_{U,t-1} - C_{U,t}) \\ &\geq \Delta_t\, I(C_t \neq C_{t-1}). \end{aligned}$$

Consequently, for every integer $n > \tilde{\sigma}$,

$$\sum_{t=\tilde{\sigma}+1}^{n} I(C_t \neq C_{t-1}) \leq \frac{\sum_{t=\tilde{\sigma}+1}^{n}(W_{t-1} - W_t)}{\min\{\Delta, \alpha\lambda W_\infty/4\}} = \frac{W_{\tilde{\sigma}} - W_n}{\min\{\Delta, \alpha\lambda W_\infty/4\}} \leq \frac{W_{\tilde{\sigma}}}{\min\{\Delta, \alpha\lambda W_\infty/4\}} < \infty.$$

Letting $n \to \infty$ shows that only finitely many hull changes occur. The hull would therefore eventually remain constant with positive width, contradicting the preceding simultaneous statement. Hence $\lim_{t\to\infty} W_t = 0$.

Finally, $q_\tau \in C_t$ and (E11) imply

$$0 \leq q_\tau - C_{L,t} \leq W_t, \quad 0 \leq C_{U,t} - q_\tau \leq W_t,$$
$$\max\{|\widetilde{L}_t - q_\tau|,\ |\widetilde{U}_t - q_\tau|\} \leq W_{t-1}, \quad 0 < r_t \leq W_{t-1}/2.$$

Therefore, almost surely on $\mathcal{G}$,

$$\tilde{\sigma} < \infty, \quad \lim_{t\to\infty} W_t = 0, \quad \lim_{t\to\infty} r_t = 0,$$
$$\lim_{t\to\infty} C_{L,t} = \lim_{t\to\infty} C_{U,t} = \lim_{t\to\infty} \widetilde{L}_t = \lim_{t\to\infty} \widetilde{U}_t = q_\tau. \tag{E14}$$

Step 2: slope calibration. For the actual query sequence, define the private secant slope b_t, weighted noise sum $\mathcal{M}_t$, and accumulated squared radii D_t:

$$b_t = \frac{\rho\{F(\widetilde{U}_t) - F(\widetilde{L}_t)\}}{2r_t}, \quad D_t := \sum_{s\leq t} r_s^2, \quad D_0 = \mathcal{M}_0 = 0,$$
$$\mathcal{M}_t = \sum_{s\leq t} r_s\{I(Z_s = 1) - a - 2r_s b_s\}.$$

Summands after termination are zero. Conditional centering gives

$$\mathbb{E}(\mathcal{M}_t - \mathcal{M}_{t-1} \mid \mathcal{F}_{t-1}) = 0,$$
$$\sum_{t\geq 1} \mathbb{E}\left\{ \left(\frac{\mathcal{M}_t - \mathcal{M}_{t-1}}{1 + D_t}\right)^2 \middle| \mathcal{F}_{t-1} \right\} \leq \sum_{t\geq 1} \frac{D_t - D_{t-1}}{(1+D_t)^2} \leq \int_0^\infty \frac{dx}{(1+x)^2} = 1.$$

The normalized martingale series converges almost surely; Kronecker's lemma (Durrett, 2019, Theorem 2.5.9) gives

$$\lim_{t\to\infty} \frac{\mathcal{M}_t}{D_t} = 0, \quad \text{almost surely on } \left\{\lim_{t\to\infty} D_t = \infty\right\}.$$

On $\mathcal{G}$, if $\lim_{t\to\infty} D_t < \infty$, then

$$\sum_{u=\tilde{\sigma}}^{\infty} W_u^2 \leq \frac{4}{\lambda^2} \sum_{s=\tilde{\sigma}+1}^{\infty} r_s^2 < \infty, \quad 0 \leq \limsup_{t\to\infty} \lfloor t/2 \rfloor W_t^2 \leq \lim_{t\to\infty} \sum_{u=\lfloor t/2 \rfloor}^{t} W_u^2 = 0.$$

Hence $\lim_{t\to\infty} tW_t^2 = 0$. In the remaining case, almost surely on $\mathcal{G}$,

$$\lim_{t\to\infty} b_t = b_*, \quad \lim_{t\to\infty} \frac{\sum_{s\leq t} r_s^2 b_s}{D_t} = b_*,$$
$$\lim_{t\to\infty} \widehat{b}_{t+1} = \lim_{t\to\infty} \max\left\{ b_{\min}, \frac{\sum_{s\leq t} r_s^2 b_s}{D_t} + \frac{\mathcal{M}_t}{2D_t} \right\} = b_*.$$

Here continuity gives the first limit, weighted averaging the second, and $b_{\min} \leq b_*$ the last.

Step 3: a fixed informative side in each phase. Define the rounded-width bound γ, phase contraction ratio η, and endpoint separation constant χ by

$$\gamma := (1+\alpha)\lambda < 1, \quad \eta := \frac{2}{3-\gamma} \in (0,1), \quad \chi := \frac{2\lambda+\gamma-1}{2\gamma} > \lambda > \frac{1}{2}.$$

Indeed, $\chi - \lambda = (2\lambda-1)(1-\gamma)/(2\gamma) > 0$. For a deterministic first report m with bounded nonempty C_{m-1}, define phase boundaries T_j and initial widths w_j by

$$T_1 = m-1, \quad w_j := W_{T_j}, \quad T_{j+1} := \inf\{t > T_j : C_t = \varnothing \text{ or } W_t \le \eta w_j\}.$$

These are stopping times. No next phase starts after an infinite boundary or empty termination; w_j is defined only for phases that start. For $T_j < t \le T_{j+1}$, nesting implies

$$\eta w_j < W_{t-1} \le w_j, \quad |\mu_t - \mu_{T_j+1}| \le \frac{w_j - W_{t-1}}{2}.$$

By rounding and $(1-\gamma)\eta/2 = 1-\eta$,

$$\begin{aligned} \widetilde{L}_t &> C_{L,t-1} + \frac{1-\gamma}{2} W_{t-1} > C_{L,T_j} + (1-\eta) w_j, \\ \widetilde{U}_t &< C_{U,t-1} - \frac{1-\gamma}{2} W_{t-1} < C_{U,T_j} - (1-\eta) w_j, \\ \frac{\lambda\eta w_j}{2} &< r_t \le \frac{\gamma w_j}{2}. \end{aligned} \tag{E15}$$

Select one factor for the entire phase, measurably at its start:

$$e_t^\star := \begin{cases} e_{t,L}, & q_\tau \ge \mu_{T_j+1}, \\ e_{t,U}, & q_\tau < \mu_{T_j+1}. \end{cases}$$

This choice is only for analysis. In the left case,

$$q_\tau - \widetilde{L}_t \ge \frac{\lambda W_{t-1} - (w_j - W_{t-1})}{2} \ge \frac{1+\lambda-1/\eta}{2} W_{t-1} = \frac{\chi\gamma}{2} W_{t-1} \ge \chi r_t;$$

the right case gives $\widetilde{U}_t - q_\tau \ge \chi r_t$. For a round $t < T_{j+1}$, $W_t > \eta w_j$ also gives

$$C_{L,t} < C_{L,T_j} + (1-\eta) w_j, \quad C_{U,t} > C_{U,T_j} - (1-\eta) w_j.$$

Since $C_{L,t} = \inf S_t$ and $C_{U,t} = \sup S_t$, there is a surviving candidate in the selected outer region. It receives the selected factor throughout the phase; endpoint attainment is unnecessary.

Step 4: conditional log drift. For a binary report $B \sim \mathrm{Bernoulli}(c-d)$, let $c \in (0,1)$ be its null mean, $d \in [0,c]$ the mean gap, and $x \ge 0$ the betting scale. With forecast $v = ce^{-x/c}$, its expected log factor is

$$\begin{aligned} &\phi(x) := \mathbb{E} \log \Lambda(B; ce^{-x/c}, c) = -\frac{(c-d)x}{c} + (1-c+d) \log \frac{1-ce^{-x/c}}{1-c}, \\ &\phi(0) = 0, \quad \phi'(0) = \frac{d}{c(1-c)}, \\ &\phi''(x) = -\frac{(1-c+d)e^{-x/c}}{c(1-ce^{-x/c})^2} \ge -\frac{1-c+d}{c(1-c)^2}. \end{aligned}$$

Taylor's formula and the range over the two outcomes imply

$$\begin{aligned} \mathbb{E} \log \Lambda(B; v, c) &\ge \frac{x}{c(1-c)} \left\{ d - \frac{x}{2} - \frac{dx}{2(1-c)} \right\}, \\ \mathrm{range}\{\log \Lambda(B; v, c)\} &= \frac{x}{c} + \log \frac{1-ce^{-x/c}}{1-c} \le \frac{x}{c(1-c)}. \end{aligned} \tag{E16}$$

The last inequality follows by integrating the derivative $1/\{c(1-ce^{-x/c})\} \le 1/\{c(1-c)\}$ from zero to x. Apply this calculation conditionally with

$$(B, c, d_t) = \begin{cases} (B_{t,L}, c_L, \rho\{\tau - F(\widetilde{L}_t)\}), & \text{left}, \\ (1-B_{t,U}, 1-c_U, \rho\{F(\widetilde{U}_t) - \tau\}), & \text{right}, \end{cases} \quad x_t := \kappa \widehat{b}_t r_t.$$

Since $\chi > 1/2$ and $0 < \kappa < 1$, choose a fixed relative calibration tolerance $\xi \in (0,1)$ such that

$$\frac{\kappa(1+\xi)}{(1-\xi)\chi} < 2.$$

Steps 1-2 imply, almost surely on $\mathcal{G} \cap \{\lim_{t\to\infty} D_t = \infty\}$,

$$\lim_{t\to\infty} \inf_{x \in C_{t-1}} f(x) = f(q_\tau), \quad \lim_{t\to\infty} \widehat{b}_t = b_*, \quad \lim_{t\to\infty} x_t = 0.$$

For all sufficiently late reports of any newly started phase,

$$d_t \geq (1-\xi) b_* \chi r_t, \quad \kappa(1-\xi) b_* r_t \leq x_t \leq \kappa(1+\xi) b_* r_t,$$
$$\frac{d_t - x_t/2 - d_t x_t/\{2(1-c)\}}{d_t} \geq 1 - \frac{\kappa(1+\xi)}{2(1-\xi)\chi} - \frac{x_t}{2(1-c)},$$
$$\liminf_{t\to\infty} \left\{ 1 - \frac{\kappa(1+\xi)}{2(1-\xi)\chi} - \frac{x_t}{2(1-c)} \right\} > 0.$$

Combining (E15) and (E16) yields fixed drift and range constants $c_1, c_2 > 0$ with

$$\mathbb{E}(\log e_t^\star \mid \mathcal{F}_{t-1}) \geq c_1 b_*^2 w_j^2, \quad \operatorname{range}(\log e_t^\star \mid \mathcal{F}_{t-1}) \leq c_2 b_* w_j. \tag{E17}$$

The finite onset depends on the path, but not on the phase partition.

Step 5: phase duration. Fix $m \geq 1$ and a concentration error probability $\delta \in (0,1)$. For phase j, let $\mathcal{M}_{j,n}$ sum the centered log factors over its first n reports, with $\mathcal{M}_{j,0} = 0$. Set subsequent increments to zero after the phase ends or q_τ is excluded, and from the first failure of (E17) at a round $t \geq m$. These predictable restrictions make $\mathcal{M}_{j,n}$ an auxiliary martingale.

Conditional on $\mathcal{F}_{T_j}$, Hoeffding's lemma and the exponential maximal inequality give, for integer $N \geq 1$ and $u \in (0,1)$,

$$\Pr\left(\max_{1\leq n\leq N} |\mathcal{M}_{j,n}| > c_2 b_* w_j \sqrt{\frac{N}{2} \log \frac{2}{u}} \,\middle|\, \mathcal{F}_{T_j} \right) \leq u.$$

Use $N = 2^k$ and $u = \delta/\{4j^2(k+1)^2\}$, where $k \geq 0$. Since

$$\sum_{j=1}^{\infty} \sum_{k=0}^{\infty} \frac{\delta}{4j^2(k+1)^2} = \frac{\delta \pi^4}{144} < \delta,$$

there is a simultaneous concentration event $\mathcal{E}_{m,\delta}$, with $\Pr(\mathcal{E}_{m,\delta}) \geq 1-\delta$, on which

$$\max_{1\leq i\leq n} |\mathcal{M}_{j,i}| \leq c_2 b_* w_j \sqrt{n \ell_{j,n}}, \quad \ell_{j,n} := \log \frac{8j^2(1+\lceil \log_2 n \rceil)^2}{\delta}, \quad j, n \geq 1.$$

Here $\ell_{j,n}$ is the simultaneous logarithmic penalty; phases that do not start impose no condition. This event does not condition on future survival.

On this event, consider the first $n \geq 2$ rounds of phase j, all occurring before the phase ends and completed by time t, with all stopping checks satisfied. Choose a candidate $q \in S_{T_j+n}$ in the selected outer region from Step 3. For $1 \leq k \leq \lfloor n/2 \rfloor$, the drift and concentration bounds imply

$$\sum_{u=T_j+k}^{T_j+n} \log e_u(q) \geq \frac{c_1 b_*^2 n w_j^2}{2} - 2 c_2 b_* w_j \sqrt{n \ell_{j,n}},$$
$$\beta_{T_j+k} \geq \frac{1}{(t+1)\{1+\log(t+1)\}^2}.$$

Iterating the capital recursion and using $\lfloor n/2 \rfloor \geq n/3$ gives

$$\alpha^{-1} > E_{T_j+n}(q) \geq A_{T_j+n}(q) \geq \sum_{k=1}^{\lfloor n/2 \rfloor} \beta_{T_j+k} \prod_{u=T_j+k}^{T_j+n} e_u(q)$$
$$\geq \frac{n \exp(c_1 b_*^2 n w_j^2/2 - 2 c_2 b_* w_j \sqrt{n \ell_{j,n}})}{3(t+1)\{1+\log(t+1)\}^2}.$$

Therefore, by $2c_2 b_* w_j \sqrt{n\ell_{j,n}} \le c_1 b_*^2 n w_j^2/4 + 4c_2^2 \ell_{j,n}/c_1$,

$$b_*^2 n w_j^2 \le C\left[\log\frac{t+1}{n} + \log\frac{3}{\alpha} + 2\log\{1+\log(t+1)\} + \ell_{j,n}\right]. \tag{E18}$$

Step 6: summation and the limiting rate. Work on $\mathcal{E}_{m,\delta}$ with bounded nonempty C_{m-1}, $q_\tau \in S_t$, and (E17) holding at every round from m through t, where $t \ge 2m-1$. Define the scaled squared widths and logarithmic penalty

$$X := (t+1)b_*^2 W_t^2, \quad u_0 := b_*^2 w_1^2, \quad \ell_t(\delta) := 1 + \log\{1+\log(t+1)\} + \log(1/\delta).$$

First suppose $X > u_0$. Let J count phases through t, including a new zero-report phase when t ends its predecessor. Then

$$W_t \le w_J \le \eta^{J-1} w_1, \quad 1 < X/u_0 \le (t+1)\eta^{2(J-1)}, \quad J \le 1 + \frac{\log(t+1)}{2|\log\eta|}.$$

Let n_j count the rounds of phase j completed by time t before the phase ends, with $n_j = 0$ if fewer than two such rounds occur. At most two reports per phase are omitted:

$$t - m + 1 \le 2J + \sum_{j=1}^{J} n_j.$$

The scale-conversion weights $a_j := W_t^2/w_j^2$ satisfy

$$a_j \le \eta^{2(J-j)}, \quad \sum_{j=1}^{J} a_j \le \frac{1}{1-\eta^2}, \quad \sum_{j=1}^{J} a_j \log(1/a_j) \le \frac{2}{e(1-\eta)}, \quad J b_*^2 W_t^2 \le \frac{u_0}{1-\eta^2}.$$

Here $x\log(1/x) \le (2/e)\sqrt{x}$ for $0 < x \le 1$ proves the third bound. For $n_j \ge 2$, (E18) reduces to

$$b_*^2 n_j w_j^2 \le C\left\{\ell_t(\delta) + \log\frac{t+1}{n_j}\right\}.$$

If $b_*^2 n_j w_j^2 > u_0$, then

$$\frac{t+1}{n_j} = \frac{X}{a_j b_*^2 n_j w_j^2} \le \frac{X}{a_j u_0}.$$

Combining the two cases, including $n_j = 0$, gives

$$b_*^2 n_j w_j^2 \le u_0 + C\{\ell_t(\delta) + \log(X/u_0) + \log(1/a_j)\}.$$

Multiply by a_j and sum:

$$\begin{aligned}\frac{X}{2} &\le (t-m+1) b_*^2 W_t^2 \\ &\le 2J b_*^2 W_t^2 + \sum_{j=1}^{J} a_j b_*^2 n_j w_j^2 \\ &\le C\{u_0 + \ell_t(\delta) + \log(X/u_0)\}.\end{aligned}$$

For $\log^+ x := \max\{0, \log x\}$, absorption gives

$$C\log(X/u_0) \le \frac{X}{4} + C\log^+(4C/u_0) \quad\Longrightarrow\quad \frac{X}{4} \le C\{u_0 + \ell_t(\delta) + \log^+(4C/u_0)\}.$$

Enlarging C yields

$$(t+1)W_t^2 \le C w_1^2 + \frac{C}{b_*^2}\left\{\ell_t(\delta) + \log^+\frac{1}{b_*^2 w_1^2}\right\}. \tag{E19}$$

The same bound holds for $X \le u_0$. It is simultaneous over all such t.

Construct $\mathcal{E}_{m,1/k}$ for every deterministic integer $m \geq 1$ and $k \geq 2$. For each fixed m,

$$\Pr\left[\left\{\bigcup_{k\geq 2} \mathcal{E}_{m,1/k}\right\}^c\right] \leq \inf_{k\geq 2} \Pr(\mathcal{E}_{m,1/k}^c) \leq \inf_{k\geq 2} \frac{1}{k} = 0.$$

This argument does not require the events to be nested. Taking a countable intersection over m therefore gives

$$\Pr\left(\bigcap_{m\geq 1} \bigcup_{k\geq 2} \mathcal{E}_{m,1/k}\right) = 1.$$

On this probability-one event and $\mathcal{G} \cap \{\lim_{t\to\infty} D_t = \infty\}$, Steps 1-4 allow an integer $m > \tilde{\sigma}$ after calibration has stabilized, and the union supplies a finite k for which $\mathcal{E}_{m,1/k}$ holds. Both choices may depend on the path, since the events were constructed for all deterministic m, k; no stopping-time assumption on m is needed. For these fixed values,

$$\lim_{t\to\infty} \frac{\ell_t(1/k)}{\log\log t} = 1, \quad \lim_{t\to\infty} \frac{b_*^2 w_1^2 + \log^+\{1/(b_*^2 w_1^2)\}}{\log\log t} = 0.$$

Thus (E19) and the finite-D_t case give

$$\limsup_{t\to\infty} \frac{b_*^2 t W_t^2}{\log\log t} \leq C, \quad \text{almost surely on } \mathcal{G}, \quad W_t = \mathcal{O}\left\{\frac{1}{\rho f(q_\tau)}\sqrt{\frac{\log\log t}{t}}\right\}.$$

Together with (E14), this proves the theorem. □

E.4 Proof of Theorem C2

Proof. Step 1: acquisition cuts. Fix a deterministic $z \in \mathcal{Z}$. For the continued Stage 1 process in Lemma E.1, define its maximal displacement

$$D_z := \sup_{t\geq 1} |\theta_t - z|.$$

Recall $B_K = \max\{c_0, 1 - c_0\}$ and the supermartingale inequality (E7). The process $V_t(z)$ satisfies

$$|\theta_t - z|^2 \leq V_t(z), \quad V_1(z) = (\theta_1 - z)^2 + \frac{B_K^2}{b_{\min}^2} \sum_{s=1}^{\infty} (s + t_0)^{-2}.$$

For every excursion threshold $r > 0$, the maximal inequality gives

$$\Pr(D_z > r) \leq \lim_{N\to\infty} \Pr\left\{\max_{1\leq t\leq N} V_t(z) > r^2\right\} \leq \min\left\{1, \frac{V_1(z)}{r^2}\right\},$$

$$\mathbb{E}D_z = \int_0^\infty \Pr(D_z > r)\, dr$$

$$\leq \int_0^{\sqrt{V_1(z)}} 1\, dr + \int_{\sqrt{V_1(z)}}^\infty \frac{V_1(z)}{r^2}\, dr = 2\sqrt{V_1(z)}.$$

The actual procedure and the continued Stage 1 process agree until acquisition or termination. Since $h_t = h_1$ during acquisition, all acquisition cut points belong to

$$\Delta\mathbb{Z} \cap (z - D_z - h_1 - \Delta,\ z + D_z + h_1 + \Delta).$$

The number of lattice points in an interval is at most its length divided by Δ plus one. Counting cuts and cells therefore gives

$$M_t \leq 1 + \left(\frac{2D_z + 2h_1 + 2\Delta}{\Delta} + 1\right) = 4 + \frac{2h_1}{\Delta} + \frac{2D_z}{\Delta} \quad \text{during acquisition.}$$

If acquisition occurs, $C_{L,\tilde{\sigma}}, C_{U,\tilde{\sigma}}$ are historical cuts, so

$$\frac{W_{\tilde{\sigma}}}{\Delta} = \frac{C_{U,\tilde{\sigma}} - C_{L,\tilde{\sigma}}}{\Delta} \le 2 + \frac{2h_1}{\Delta} + \frac{2D_z}{\Delta}.$$

Step 2: the uniform cell bound. The mesh bound in (E12) uses only the dyadic rounding rule and the monotonicity of W_t, so it holds for any fixed $\alpha, \lambda \in (0,1)$. Surviving cells have disjoint interiors and length at least the current mesh. Thus $M_{\tilde{\sigma}} \le W_{\tilde{\sigma}}/\Delta$, and (E12) gives, for $t > \tilde{\sigma}$ before termination,

$$\begin{aligned} M_t \le \frac{W_t}{\Delta_t} &\le \frac{W_{t-1}}{\min\{\Delta, \alpha\lambda W_{t-1}/4\}} \\ &= \max\left\{\frac{W_{t-1}}{\Delta}, \frac{4}{\alpha\lambda}\right\} \le \max\left\{\frac{W_{\tilde{\sigma}}}{\Delta}, \frac{4}{\alpha\lambda}\right\}. \end{aligned}$$

After termination $M_t = 0$. With the deterministic counting constant $K := \max\{4 + 2h_1/\Delta, 4/(\alpha\lambda)\}$, both stages imply

$$\sup_{t\ge 0} M_t \le K + \frac{2D_z}{\Delta}.$$

Consequently, for $\delta \in (0,1)$,

$$\begin{aligned} \Pr\left\{\sup_{t\ge 0} M_t > K + \frac{2}{\Delta}\sqrt{\frac{V_1(z)}{\delta}}\right\} &\le \delta, \\ \mathbb{E}\left(\sup_{t\ge 0} M_t\right) &\le K + \frac{4}{\Delta}\sqrt{V_1(z)} < \infty. \end{aligned}$$

In particular, $\Pr(\sup_{t\ge 0} M_t < \infty) = 1$. No acquisition or density assumption was used.

Step 3: exactness and storage. Let $(x_-, x_+]$ be an individual surviving cell with common active capital A before report t. This is equivalent to storing total capital $R_{t-1} + A$. An interior cut x makes the replacement

$$((x_-, x_+], A) \longmapsto \big\{((x_-, x], A),\ ((x, x_+], A)\big\}.$$

After splitting at $\widetilde{L}_t$ and $\widetilde{U}_t$, all candidates within each child cell J receive the same betting factor. Hence, for $q, q' \in J$,

$$e_t(q) = e_t(q'), \quad A_t(q) = A_t(q') = (A + \beta_t)e_t(q) =: A',$$

where A' is its updated active capital. The update retains exactly

$$\{q \in J : E_t(q) < \alpha^{-1}\} = \begin{cases} J, & R_t + A' < \alpha^{-1}, \\ \varnothing, & R_t + A' \ge \alpha^{-1}. \end{cases}$$

Starting from $S_0 = \mathbb{R}$ and $A_0(q) = 0$, induction using

$$S_t = S_{t-1} \cap \{q : E_t(q) < \alpha^{-1}\}, \quad C_t = [\inf S_t, \sup S_t] \quad (S_t \ne \varnothing)$$

shows that the cell representation exactly reproduces the candidate-wise updates for the rounded query sequence. Existing boundary cuts change nothing; deleted gaps remain deleted. Infinite endpoints use the natural interval conventions.

Each surviving cell stores two endpoints and one capital value, and splitting creates at most two additional records. The slope sums and tracker, query, lattice, and reserve states require only a fixed number of stored values. Thus, the peak working memory satisfies

$$\text{storage}_{\text{peak}} = \mathcal{O}\left(1 + \sup_{t\ge 0} M_t\right).$$

By Step 2, this bound is almost surely finite and has finite expectation. □