Supplementary material

EC.1. Proof of Proposition 1

Consider a two-dimensional pricing setting $d=2$, with feature map $\phi(x,p)=(0.5x,0.5p)$ for $x\in[-1,1]$, $p\in[0,1]$ and the hypothetical demand model $\theta^*=(0.5,-0.5)$. The logistic demand model $\Pr[d_t=1|p_t,x_t]=\frac{e^{\zeta\phi(x_t,p_t)^\top\theta^*}}{1+e^{\zeta\phi(x_t,p_t)^\top\theta^*}}$ is used, with $\zeta=8$, or equivalently $\Pr[d_t=1|p_t,x_t]=\frac{e^{2x_t-2p_t}}{1+e^{2x_t-2p_t}}$. It is easy to verify that, with $x_t=1$, the optimal price is $p_t=1$; with $x_t=-1$, the optimal price is $p_t\approx 0.524$. Let $r_1^*=\max_{p\in[0,1]}p_t\Pr[d_t=1|p_t,x_t=1]$, $r_{-1}^*=\max_{p\in[0,1]}p_t\Pr[d_t=1|p_t,x_t=-1]$ be the maximum expected revenue when priced optimally for $x_t=1$ and $x_t=-1$. Define

$$\Delta_r^*:=\min_{p\in[0,1]}\max_{j\in\{1,-1\}}\left\{r_j^*-p\times\Pr[d=1|p,x=j]\right\}$$

as the minimum regret of any fixed price $p\in[0,1]$ under either context $x=1$ or $x=-1$. It is easy to verify that $\Delta_r^*\geq 0.001>0$ and therefore $\Delta_r^*=\Omega(1)$.

For each time t, let $\mathcal{H}_{t-1}:=\{x_\tau,y_\tau,p_\tau\}_{\tau<t}$ be the complete history prior to time t. Because the policy π must be non-anticipating and satisfies the (ε,δ)-differential privacy in Definition 1, for any $p\in[0,1]$ and $x,x'\in[-1,1]$, it holds that

$$\Pr[p_t=p|\mathcal{H}_{t-1},x_t=x]\leq e^\varepsilon\Pr[p_t=p|\mathcal{H}_{t-1},x_t=x']+\delta \tag{EC.1}$$

Let the environment be such that $x_t=\pm1$ uniformly at random. Then the expected regret of a policy π satisfying Eq. (EC.1) at time t can be lower bounded by

$$\begin{aligned}
&\frac{1}{2}\sum_{j\in\{1,-1\}}\int_0^1(r_j^*-p\times\Pr[d=1|p,x_t=j])\times\Pr[p_t=p|\mathcal{H}_{t-1},x_t=j]\mathrm{d}p\\
&\geq\frac{1}{2}\Delta_r^*\times\int_0^1\min_{j\in\{1,-1\}}\Pr[p_t=p|\mathcal{H}_{t-1},x_t=j]\mathrm{d}p\\
&\geq\frac{1}{2}\Delta_r^*\times\int_0^1\frac{\max\{0,\Pr[p_t=p|\mathcal{H}_{t-1},x_t=1]-\delta\}}{e^\varepsilon}\mathrm{d}p &\text{(EC.2)}\\
&\geq\frac{1}{4}\Delta_r^*\times\int_0^1\max\{0,\Pr[p_t=p|\mathcal{H}_{t-1},x_t=1]-1/4\}\mathrm{d}p\\
&\geq\frac{1}{4}\Delta_r^*\times(1-1/4) &\text{(EC.3)}\\
&\geq\frac{1}{4}\times\frac{3}{4}\Delta_r^*\geq\frac{3}{16}\Delta_r^*=\Omega(1),
\end{aligned}$$

where in Eq. (EC.2) we apply Eq. (EC.1), and Eq. (EC.3) holds because $\int_0^1 \Pr[p_t = p|\mathcal{H}_{t-1}, x_t = 1]\mathrm{d}p = 1$. Summing over all T time periods we obtain the desired $\Omega(T)$ lower bound on the policy π.

EC.2. Proof of Proposition 2

Proof of Proposition 2. Let D, D' be a pair of neighboring databases differing only on time t. For every time period τ, suppose $p_\tau = f_\tau(x_\tau, a_1, \cdots, a_{\tau-1})$ for some deterministic function f_τ. Let $\mathcal{P}_{>t}$ be a measurable set over $p_{t+1}, \cdots, p_T$. Define $\mathcal{A} := \{a_1, \cdots, a_T : f_\tau(x_\tau, a_1, \cdots, a_{\tau-1}) = p_\tau, \forall \tau > t\}$. Because $x_\tau = x'_\tau$ for all $\tau > t$, the set $\mathcal{A}$ is the same under D and D'. Note that our definition of anticipating privacy plays a key role here since we only care about the future prices, while D and D' only differs at time t. Therefore, $\Pr[p_{t+1}, \cdots, p_T \in \mathcal{P}_{>t}|\pi, D] = \Pr[a_1, \cdots, a_T \in \mathcal{A}|\pi, D]$ and $\Pr[p_{t+1}, \cdots, p_T \in \mathcal{P}_{>t}|\pi, D'] = \Pr[a_1, \cdots, a_T \in \mathcal{A}|\pi, D']$. Additionally, because $(a_1, \cdots, a_T)$ are (ε, δ)-differentially private it holds that $\Pr[a_1, \cdots, a_T \in \mathcal{A}|\pi, D] \le e^\varepsilon \Pr[a_1, \cdots, a_T \in \mathcal{A}|\pi, D'] + \delta$. Subsequently, $\Pr[p_{t+1}, \cdots, p_T \in \mathcal{P}_{>t}|\pi, D] \le e^\varepsilon \Pr[p_{t+1}, \cdots, p_T \in \mathcal{P}_{>t}|\pi, D'] + \delta$, which is to be demonstrated. □

EC.3. Proofs of technical lemmas in Section 6.1

EC.3.1. Proof of Proposition 3.

Because $\|\phi_t\|_2 \le 1$ almost surely, Section 3 and Algorithm 1 in (Dwork et al. 2014) shows that each $\widehat{\Sigma}(\ell)$ satisfies (ε', δ')-differential privacy. Additionally, Theorem 3.5 in (Chan et al. 2011) shows that each Σ_n^p involves at most $m = \lceil \log_2 T \rceil$ differentially private partial sums. By switching the basic composition argument in Theorem 3.5 of (Chan et al. 2011) to advanced composition (Corollary 1), we have that the entire procedure satisfies (ε, δ)-differential privacy. □

EC.3.2. Proof of Lemma 1.

First fix n. Invoking Corollary 2.3.6 from Tao (2012), there exist constants $c, C > 0$ such that, for every $a \ge C$,

$$\Pr[\|W^n\|_{\mathrm{op}} > \sigma_{\varepsilon',\delta'} a\sqrt{d}] \le Ce^{-cad}.$$

Equating the right-hand side of the above inequality with $1/T^2$ and applying union bound over all T periods, we have with probability $1 - O(T^{-1})$ that

$$\|W^n\|_{\mathrm{op}} \le \sigma_{\varepsilon',\delta'} \min\big\{C\sqrt{d}, \frac{\ln(CT^2)}{c\sqrt{d}}\big\}, \quad \forall n.$$

Note that for each n, the difference between Σ_n^p and Σ_n involves at most $m = \lceil \log_2 T \rceil$ noise matrices W. Subsequently, for every n,

$$\begin{aligned}
\|\Sigma_n^p - \Sigma_n\|_{\mathrm{op}} &\le m \max_n \|W^n\|_{\mathrm{op}} \le \sigma_{\varepsilon',\delta'} \ln T \times O(\sqrt{d}, \ln(T)/\sqrt{d}) \le \sigma_{\varepsilon',\delta'} \times O(\sqrt{d}\ln^2 T) \\
&\le \frac{2\lceil \log_2 T\rceil \ln(1/\delta')}{\varepsilon} \times \sqrt{2\ln(1.25/\delta')} \times O(\sqrt{d}\ln^2 T) = O(\varepsilon^{-1}\sqrt{d}\ln^{4.5}(T/\delta)),
\end{aligned}$$

which is to be demonstrated. □

EC.3.3. Proof of Corollary 2.

Let $\{\lambda_j^n, \widetilde{\lambda}_j^n\}_{j=1}^d$ be the eigenvalues of Λ_n, Λ_n^p, respectively, which are all real since Σ_n, Σ_n^p are symmetric. By Lemma 1 and Weyl's theorem, we have with probability $1 - O(T^{-1})$ uniformly over all n that

$$\left|\lambda_j^n - \widetilde{\lambda}_j^n\right| = O(\varepsilon^{-1}\sqrt{d}\ln^{4.5}(T/\delta)), \quad \forall j. \tag{EC.4}$$

On the other hand, because $\Sigma_n = \sum_{t\le n} \phi_t\phi_t^\top$ is positive semi-definite, $\lambda_j^n \ge \rho \ge \varepsilon^{-1} d\sqrt{d}\ln^5(T/\delta)$ for all j and n. Subsequently, there exists a universal constant $C_T < \infty$ such that for any $T \ge C_T$, with probability $1 - O(T^{-1})$ uniformly over all n and j that

$$\left|\lambda_j^n - \widetilde{\lambda}_j^n\right| \le \frac{0.1}{d}\lambda_j^n.$$

Subsequently, $\det(\Lambda_n^p) \le (1 + \frac{0.1}{d})^d \det(\Lambda_n) \le e^{0.1}\det(\Lambda_n) \le 1.11\det(\Lambda_n)$ and $\det(\Lambda_n^p) \ge (1 - \frac{0.1}{d})^d \det(\Lambda_n) \ge e^{-0.1}\det(\Lambda_n) \ge 0.9\det(\Lambda_n)$, which are to be demonstrated. □

EC.4. Proofs of technical lemmas in Section 6.2

EC.4.1. Proof of Proposition 4.

The proof of privacy amounts to verify the several conditions in Theorem 2 of Kifer et al. (2012) hold true.

First we verify that the objective $\sum_{t<n} -\ln p(y_t|\phi_t, \theta)$ is convex with respect to θ. Recall that $p(y_t|\phi_t,\theta) = \exp\{\zeta(y_t\phi_t^\top\theta - m(\phi_t^\top\theta)) + h(y_t,\zeta)\}$. Taking the first and the second derivatives, we have that

$$\begin{aligned}
-\nabla_\theta \ln p(y_t|\phi_t,\theta) &= \zeta(m'(\phi_t^\top\theta) - y_t)\phi_t = \zeta(f(\phi_t^\top\theta) - y_t)\phi_t; \\
-\nabla_\theta^2 \ln p(y_t|\phi_t,\theta) &= \zeta f'(\phi_t^\top\theta)\phi_t\phi_t^\top.
\end{aligned}$$

Because $\zeta \ge G^{-1} > 0$ and $f'(\cdot) \ge K^{-1} > 0$, we have that $-\nabla_\theta^2 \ln p(y_t|\phi_t,\theta) \succeq 0$ for all θ. This shows that the objective function is convex with respect to θ.

We next upper bound the gradients and Hessian matrices. Recall that we assume the realized demands y_t are uniformly bounded by B_Y in all periods and databases. Furthermore, $\zeta \leq G$, $f(\cdot) \in [0,1]$, $f'(\cdot) \leq K$ and $\|\phi_t\|_2 \leq 1$ almost surely. Hence, $\|-\nabla_\theta \ln p(y_t|\phi_t,\theta)\|_2 \leq (B_Y+1)G = B_1$, $\|-\nabla^2_\theta \ln p(y_t|\phi_t,\theta)\|_{\mathrm{op}} \leq KG = B_2$ for all t and θ. Invoking Theorem 2 of (Kifer et al. 2012) we complete the proof of Proposition 4. $\square$

EC.4.2. Proof of Lemma 2

To prove Lemma 2, we need to introduce two important technical lemmas as follows. The proofs of these two technical lemmas will be deferred to the end of this section.

LEMMA EC.1. *Fix $n \in \{1,2,\cdots,T\}$. For every $t<n$, define $\ell_t(\theta) = \mathbb{E}_{y\sim p(\cdot|\phi_t,\theta^*)}[-\ln p(y|\phi_t,\theta)]$. Define $F_n(\theta) = L_n(\theta) + \frac{\rho}{2}\|\theta\|_2^2 + w^\top\theta$ where $L_n(\theta) = \sum_{t<n} \ell_t(\theta)$ and $w \sim \mathcal{N}(0, \nu^2_{\varepsilon,\delta} I)$. Let θ^*_ρ be defined as*

$$\theta^*_\rho = \arg\min_{\theta\in\mathbb{R}^d} F_n(\theta).$$

If $\rho \geq 5\nu_{\varepsilon,\delta}\sqrt{5d\ln T}$ then with probability $1-O(T^{-2})$ it holds that $\|\theta^*_\rho\|_2 \leq 1.5$.

LEMMA EC.2. *Fix $n \in \{1,2,\cdots,T\}$. For every $t<n$, define $\widehat{\ell}_t(\theta) = -\ln p(y_t|\phi_t,\theta)$. Define $\widehat{F}_n(\theta) = \widehat{L}_n(\theta) + \frac{\rho}{2}\|\theta\|_2^2 + w^\top\theta$ where $\widehat{L}_n(\theta) = \sum_{t<n}\widehat{\ell}_t(\theta)$ and $w \sim \mathcal{N}(0,\nu^2_{\varepsilon,\delta} I)$. Let $\widehat{\theta}_\rho$ be defined as*

$$\widehat{\theta}_\rho = \arg\min_{\theta\in\mathbb{R}^d} \widehat{F}_n(\theta).$$

If $\rho \geq \max\{5\nu_{\varepsilon,\delta}\sqrt{5d\ln T}, 2+48s^2G^2Kd\ln T\}$ then with probability $1-O(T^{-2})$ it holds that $\|\widehat{\theta}_\rho\|_2 < 2$.

At a higher level, Lemmas EC.1 and EC.2 show that, if the smoothing term $\frac{\rho}{2}\|\theta\|_2^2$ is not too small, both global minimizers of F_n and $\widehat{F}_n$ have bounded norms with high probability. Despite their apparent similarities, the proof strategies of Lemmas EC.1 and EC.2 are quite different. Lemma EC.1 is proved by noting that θ^* is the global minimizer of $L_n(\cdot)$, and any θ^*_ρ with norm substantially larger than $\|\theta^*\|_2 \leq 1$ must incur a large penalty through the $\frac{\rho}{2}\|\theta\|_2^2$ term when ρ is not too small. On the other hand, Lemma EC.2 is proved by noting that $\widehat{F}_n$ is strongly convex and $|\widehat{F}_n - F_n|$ is small, and therefore $\widehat{\theta}_\rho$ cannot deviate too much from θ^*_ρ.

With Lemmas EC.1 and EC.2 in place, we are ready to prove Lemma 2. First, by Lemma EC.2 we know that with probability $1-O(T^{-2})$, $\|\widehat{\theta}^p_n\|_2 < 2$. The remainder of this proof is conditioned on the event that $\|\widehat{\theta}^p_n\|_2 < 2$.

Because $\|\widehat{\theta}^p_n\|_2 < 2$, the constraints in the optimization problem

$$\widehat{\theta}^p_n = \arg\min_{\|\theta\|_2\leq 2} \widehat{L}_n(\theta) + \frac{\rho}{2}\|\theta\|_2^2 + w^\top\theta$$

are not active. Therefore, by first-order KKT conditions we have that

$$0 = \nabla \widehat{L}_n(\widehat{\theta}_n^p) + \rho\widehat{\theta}_n^p + w = \sum_{t<n} \zeta(f(\phi_t^\top \widehat{\theta}_n^p) - y_t)\phi_t + \rho\widehat{\theta}_n^p + w. \tag{EC.5}$$

Recall the definition that $\xi_t = y_t - f(\phi_t^\top \theta^*)$. Re-arranging terms in Eq. (EC.5) we obtain

$$\sum_{t<n} (f(\phi_t^\top \widehat{\theta}_n^p) - f(\phi_t^\top \theta^*))\phi_t = \sum_{t<n} \xi_t\phi_t - \zeta^{-1}(\rho\widehat{\theta}_n^p + w).$$

By the mean-value theorem, there exists $\widetilde{\theta} = \theta^* + \lambda(\widehat{\theta}_n^p - \theta^*)$ for some $\lambda \in [0,1]$ such that $\sum_{t<n}(f(\phi_t^\top \widehat{\theta}_n^p) - f(\phi_t^\top \theta^*))\phi_t = \sum_{t<n} f'(\phi_t^\top \widetilde{\theta})(\widehat{\theta}_n^p - \theta^*)^\top \phi_t\phi_t$. Subsequently,

$$\sum_{t<n} f'(\phi_t^\top \widetilde{\theta})(\widehat{\theta}_n^p - \theta^*)^\top \phi_t\phi_t^\top = \sum_{t<n} \xi_t\phi_t^\top - \zeta^{-1}(\rho[\widehat{\theta}_n^p]^\top + w^\top).$$

Multiplying both sides of the above equality by $(\widehat{\theta}_n^p - \theta^*)$ and noting that $\zeta^{-1} \le G$, $\|\widehat{\theta}_n^p\|_2 < 2$, we have

$$(\widehat{\theta}_n^p - \theta^*)^\top \left[\sum_{t<n} f'(\phi_t^\top \widetilde{\theta})\phi_t\phi_t^\top\right](\widehat{\theta}_n^p - \theta^*) \le \left|\sum_{t<n} \xi_t\phi_t^\top(\widehat{\theta}_n^p - \theta^*)\right| + G(2\rho + \|w\|_2)\|\widehat{\theta}_n^p - \theta^*\|_2.$$

Note that, because $\|\theta^*\|_2 \le 1$ and $\|\widehat{\theta}_n^p\|_2 < 2$, we have $\|\widetilde{\theta}\|_2 \le 2$. Subsequently, $f'(\phi_t^\top \widetilde{\theta}) \ge K^{-1}$ for all t. Invoking also Lemma EC.3 and Eq. (EC.9), we have with probability $1 - O(T^{-2})$ that

$$\frac{1}{K}(\widehat{\theta}_n^p - \theta^*)^\top \Sigma_n(\widehat{\theta}_n^p - \theta^*) \le s\sqrt{3d\ln T}\sqrt{(\widehat{\theta}_n^p - \theta^*)^\top \Sigma_n(\widehat{\theta}_n^p - \theta^*)} + G(2\rho + \nu_{\varepsilon,\delta}\sqrt{5d\ln T})\|\widehat{\theta}_n^p - \theta^*\|_2.$$

For notational simplicity define $\Delta_\rho^2 = (\widehat{\theta}_n^p - \theta^*)^\top \Lambda_n(\widehat{\theta}_n^p - \theta^*)$ where $\Lambda_n = \Sigma_n + \rho I$. Because $\|\theta^*\|_2 \le 1$, $\|\widehat{\theta}_n^p\|_2 < 2$ we have $\|\widehat{\theta}_n^p - \theta^*\|_2 \le 3$. Subsequently,

$$\Delta_\rho^2 \le sK\sqrt{3d\ln T}\Delta_\rho + [G(2\rho + \nu_{\varepsilon,\delta}\sqrt{5d\ln T}) + 3\rho]\|\widehat{\theta}_n^p - \theta^*\|_2. \tag{EC.6}$$

Dividing both sides of Eq. (EC.6) by Δ_ρ and noting that $\Delta_\rho \ge \sqrt{\rho}\|\widehat{\theta}_n^p - \theta^*\|_2$ with $\rho \ge 1$, we obtain with probability $1 - O(T^{-2})$ that

$$\Delta_\rho \le sK\sqrt{3d\ln T} + (2G+3)\sqrt{\rho} + G\nu_{\varepsilon,\delta}\sqrt{5d\ln T}. \tag{EC.7}$$

The first inequality in Lemma 2 is thus proved.

We next prove the sharpened inequality with the additional condition that $\lambda_{\min}(\Sigma_n) \ge \lambda_0 = [\frac{(2G+3)\rho}{\sqrt{5d\ln T}} + \nu_{\varepsilon,\delta}G]^2$. With this condition, we have $\Delta_\rho \ge \sqrt{\lambda_0 + \rho}\|\widehat{\theta}_n^p - \theta^*\|_2 \ge [\frac{(2G+3)\rho}{\sqrt{5d\ln T}} + \nu_{\varepsilon,\delta}G]\|\widehat{\theta}_n^p - \theta^*\|_2$.

Subsequently, dividing both sides of Eq. (EC.6) by Δ_ρ we obtain

$$\Delta_\rho \le sK\sqrt{3d\ln T} + \sqrt{5d\ln T} \le 4sK\sqrt{d\ln T},$$

which is to be demonstrated.

EC.4.2.1. Proof of Lemma EC.1 Because $\mathbb{E}[y_t|\phi_t,\theta^*] = f(\phi_t^\top\theta^*)$ and $-\nabla\ln p(y_t|\phi_t,\theta) = \zeta(f(\phi_t^\top\theta) - y_t)\phi_t$, we have that $\nabla\ell_t(\theta) = -\nabla\mathbb{E}[\ln p(y|\phi_t,\theta)|\theta^*] = \zeta(f(\phi_t^\top\theta) - f(\phi_t^\top\theta^*))\phi_t$. It is easy to verify that $L_n = \sum_{t<n}\ell_t$ is convex and $\nabla L_n(\theta^*) = 0$. Hence, θ^* is the global minimizer of L_n and therefore $L_n(\theta^*) \le L_n(\theta_\rho^*)$.

Assume by way of contradiction that $\|\theta_\rho^*\|_2 > 1.5$. Because $\|\theta^*\|_2 \le 1$ and $L_n(\theta^*) \le L_n(\theta_\rho^*)$, we have that

$$\begin{aligned}
F_n(\theta_\rho^*) - F_n(\theta^*) &\ge \frac{\rho}{2}\|\theta_\rho^*\|_2^2 - \frac{\rho}{2}\|\theta^*\|_2^2 - \left|\langle w, \theta_\rho^* - \theta^*\rangle\right| \ge \frac{\rho}{2}\|\theta_\rho^*\|_2^2 - \frac{\rho}{2} - \|w\|_2(\|\theta^*\|_2 + \|\theta_\rho^*\|_2)\\
&\ge \frac{\rho}{2}\|\theta_\rho^*\|_2^2 - \frac{\rho}{2} - \|w\|_2 - \|w\|_2\|\theta_\rho^*\|_2\\
&= \left(\frac{\rho}{2}\|\theta_\rho^*\|_2 - \|w\|_2\right)\|\theta_\rho^*\|_2 - \left(\frac{\rho}{2} + \|w\|_2\right).
\end{aligned} \tag{EC.8}$$

Recall that $w \sim \mathcal{N}(0, \nu_{\varepsilon,\delta}^2 I)$. Hence, $\|w\|_2^2/\nu_{\varepsilon,\delta}^2$ follows χ_d^2 distribution. Invoking concentration inequalities of χ^2 random variables from (Laurent & Massart 2000), we have with probability $1-u$ that $\|z\|_2^2 \le d + 2\sqrt{d\ln(1/u)}$ for $z \sim \mathcal{N}(0, I_d)$. Subsequently, with probability $1 - O(T^{-2})$ it holds that

$$\|w\|_2^2 \le \nu_{\varepsilon,\delta}^2(d + 4\sqrt{d\ln T}) \le 5\nu_{\varepsilon,\delta}^2 d\ln T. \tag{EC.9}$$

With Eq. (EC.9) and the condition that $\rho \ge 5\nu_{\varepsilon,\delta}\sqrt{5d\ln T}$, we have that $\rho \ge 5\|w\|_2$ with probability $1 - O(T^{-2})$. Subsequently, with probability $1 - O(T^{-2})$, Eq. (EC.8) can be simplified to

$$F_n(\theta_\rho^*) - F_n(\theta^*) \ge \left(\frac{\rho}{2}\frac{3}{2} - \frac{\rho}{5}\right)\frac{3}{2} - \frac{\rho}{2} - \frac{\rho}{5} = \frac{\rho}{8} > 0,$$

which contradicts the definition that θ_ρ^* is the global minimizer of F_n.

EC.4.2.2. Proof of Lemma EC.2 To prove Lemma EC.2 we first establish the following technical lemma, which will also be useful in later proofs.

LEMMA EC.3. *Let $\xi_t = y_t - f(\phi_t^\top\theta^*)$, which are centered sub-Gaussian random variables with sub-Gaussian parameter s^2. Let $\Sigma_n = \sum_{t<n}\phi_t\phi_t^\top$. Then with probability $1 - O(T^{-2})$, it holds uni-*

formly over all $\psi\in\mathbb{R}^d$ *that*

$$\Big|\sum_{t<n}\xi_t\phi_t^\top\psi\Big|\leq s\sqrt{3d\ln T}\times\sqrt{\psi^\top(\Sigma_n+I_d)\psi}.$$

Proof of Lemma EC.3. Abbreviate $\widetilde\Lambda_n=\Sigma_n+I_d$. Invoking Theorem 1 of (Abbasi-Yadkori et al. 2011), we have with probability $1-u$ that

$$\begin{aligned}\Big\|\sum_{t<n}\xi_t\phi_t\Big\|_{\widetilde\Lambda_n^{-1}}^2&=\Big(\sum_{t<n}\xi_t\phi_t\Big)^\top(\Sigma_n+I_d)^{-1}\Big(\sum_{t<n}\xi_t\phi_t\Big)\leq 2s^2\ln\left(\frac{\det(\widetilde\Lambda_n)^{1/2}\det(I_d)^{1/2}}{u}\right)\\&\overset{(a)}{\leq}2s^2\ln\left(\frac{n^{d/2}}{u}\right)\leq s^2d\ln(n/u).\end{aligned}$$

Here Eq. (a) holds because $\|\phi_t\|_2\leq 1$ almost surely, and hence $\|\Sigma_n\|_{\mathrm{op}}\leq n-1$. Taking $u=1/T^2$ we have with probability $1-O(T^{-2})$ that

$$\Big\|\sum_{t<n}\xi_t\phi_t\Big\|_{\widetilde\Lambda_n^{-1}}\leq s\sqrt{3d\ln T}.$$

Subsequently, by Cauchy-Schwarz inequality we have

$$\Big|\sum_{t<n}\xi_t\phi_t^\top\psi\Big|\leq\Big\|\sum_{t<n}\xi_t\phi_t\Big\|_{\widetilde\Lambda_n^{-1}}\|\psi\|_{\widetilde\Lambda_n}\leq s\sqrt{3d\ln T}\times\sqrt{\psi^\top(\Sigma_n+I_d)\psi},$$

which is to be demonstrated. □

We now return to the proof of Lemma EC.2. By Lemma EC.1, with probability $1-O(T^{-2})$ it holds that $\|\theta_\rho^*\|_2\leq 1.5$, where θ_ρ^* is the global minimizer of F_n. The rest of the proof is conditioned on the event that $\|\theta_\rho^*\|_2\leq 1.5$.

The Hessian matrices of F_n and $\widehat F_n$ can be spectrally lower bounded as

$$\nabla^2\widehat F_n(\theta)=\nabla^2F_n(\theta)=\sum_{t<n}\zeta f'(\phi_t^\top\theta)\phi_t\phi_t^\top+\rho I\succeq\frac{1}{GK}\Sigma_n+\rho I,\quad\forall\|\theta\|_2\leq 2,\tag{EC.10}$$

where $\Sigma_n=\sum_{t<n}\phi_t\phi_t^\top$ and $f'(z)\geq K^{-1}$ for all $|z|\leq 2$. Subsequently, for any $\|\theta\|_2\leq 2$ it holds that

$$\begin{aligned}F_n(\theta)&\geq F_n(\theta_\rho^*)+\frac12(\theta-\theta_\rho^*)^\top\left(\frac{1}{GK}\Sigma_n+\rho I\right)(\theta-\theta_\rho^*)\\&=F_n(\theta_\rho^*)+\frac{1}{2GK}(\theta-\theta_\rho^*)\Sigma_n(\theta-\theta_\rho^*)+\frac{\rho}{2}\|\theta-\theta_\rho^*\|_2^2,\end{aligned}\tag{EC.11}$$

where the first inequality holds because θ_ρ^* is the global minimizer of F_n and therefore $\nabla F_n(\theta_\rho^*)=0$.

On the other hand, recall the definitions of $F_n, \widehat{F}_n$ as follows:

$$F_n(\theta) = \sum_{t<n} \zeta(m(\phi_t^\top \theta) - f(\phi_t^\top \theta^*)\phi_t^\top \theta) - \mathbb{E}_{\theta^*}[h(y,\zeta)] + \frac{\rho}{2}\|\theta\|_2^2 + w^\top \theta; \tag{EC.12}$$

$$\widehat{F}_n(\theta) = \sum_{t<n} \zeta(m(\phi_t^\top \theta) - y_t\phi_t^\top \theta) - h(y_t,\zeta) + \frac{\rho}{2}\|\theta\|_2^2 + w^\top \theta. \tag{EC.13}$$

Comparing Eqs. (EC.12,EC.13) and noting the definition that $y_t = f(\phi_t^\top \theta^*) + \xi_t$, we have for every $\theta \in \mathbb{R}^d$ that

$$\left|[\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*)] - [F_n(\theta) - F_n(\theta_\rho^*)]\right| = \left|\sum_{t<n} \zeta\xi_t\phi_t^\top(\theta - \theta_\rho^*)\right|.$$

Invoking Lemma EC.3 we have with probability $1 - O(T^{-2})$ for all $\theta \in \mathbb{R}^d$ that

$$\left|[\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*)] - [F_n(\theta) - F_n(\theta_\rho^*)]\right| \leq sG\sqrt{3d\ln T}\sqrt{(\theta - \theta_\rho^*)^\top(\Sigma_n + I)(\theta - \theta_\rho^*)}. \tag{EC.14}$$

Define $\Delta(\theta)^2 := (\theta - \theta_\rho^*)^\top(\Sigma_n + I)(\theta - \theta_\rho^*)$. Because $G, K \geq 1$, we have that $(\theta - \theta_\rho^*)^\top\Sigma_n(\theta - \theta_\rho^*) \geq \Delta(\theta)^2 - \frac{1}{2}\|\theta - \theta_\rho^*\|_2^2$. Subsequently, combining Eqs. (EC.11,EC.14), we have with probability $1 - O(T^{-2})$ for all $\|\theta\|_2 \leq 2$ that

$$\begin{aligned}\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*) &\geq \frac{\Delta(\theta)^2}{2GK} + \frac{\rho - 1}{2}\|\theta - \theta_\rho^*\|_2^2 - sG\sqrt{3d\ln T}\Delta(\theta) \\ &= \Delta(\theta)\left(\frac{\Delta(\theta)}{2GK} - sG\sqrt{3d\ln T}\right) + \frac{\rho - 1}{2}\|\theta - \theta_\rho^*\|_2^2. \end{aligned} \tag{EC.15}$$

Now consider any θ such that $\|\theta - \theta_\rho^*\|_2 = 0.5$. Because $\|\theta_\rho^*\|_2 \leq 1.5$, we have $\|\theta\|_2 \leq 2$. We will lower bound the right-hand side of Eq. (EC.15) by a case analysis:

1. *Case 1:* $\Delta(\theta) > 2sG^2K\sqrt{3d\ln T}$. In this case, the first term in Eq. (EC.15) is strictly positive because $\Delta(\theta) \geq \|\theta - \theta_\rho^*\|_2 \geq 1/2$. Therefore, $\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*) > 0$;
2. *Case 2:* $\Delta(\theta) \leq 2sG^2K\sqrt{3d\ln T}$. In this case, because $\rho \geq 1$, Eq. (EC.15) can be simplified to

$$\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*) \geq \frac{\rho - 1}{2}\|\theta - \theta_\rho^*\|_2^2 - \Delta(\theta) \times sG\sqrt{3d\ln T} \geq \frac{\rho - 1}{2}\frac{1}{4} - 6s^2G^2Kd\ln T.$$

   Under the condition $\rho \geq 2 + 48s^2G^2Kd\ln T$, the right-hand side of the above inequality is strictly positive. Therefore, $\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*) > 0$.

Combining the above two cases, we have proved that with probability $1 - O(T^{-2})$, for every $\|\theta - \theta_\rho^*\|_2 = 0.5$, $\widehat{F}_n(\theta) - \widehat{F}_n(\theta_\rho^*) > 0$. Since F_n is convex, this means that the global minimizer of $\widehat{F}_n$ must be contained on the interior of $\{\theta : \|\theta - \theta_\rho^*\|_2 \leq 0.5\} \subseteq \{\theta : \|\theta\|_2 \leq 2\}$. This completes the proof of Lemma EC.2.

EC.5. Proof of Theorem 1

Our first lemma shows the necessary and sufficient conditions for the $\det(\Lambda_n^p) \geq 2\det(\Lambda^p)$ condition in Step 7 of Algorithm 1 to be met:

LEMMA EC.4. *For every time period n recall the definition that $\Sigma_n = \sum_{t<n} \phi_t\phi_t^\top$ and $\Lambda_n = \Sigma_n + \rho I$. Let also Λ^p be the "current" private sample covariance maintained in Algorithm 1, and Λ be the non-private version of the sample covariance corresponding to Λ^p. With probability $1 - O(T^{-1})$ the following holds for every time period n:*

1. *If $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ then $\det(\Lambda_n) \geq 1.5\det(\Lambda)$;*
2. *If $\det(\Lambda_n) > 2.5\det(\Lambda)$ then $\det(\Lambda_n^p) > 2\det(\Lambda^p)$.*

Proof of Lemma EC.4. Obvious by invoking Corollary 2. More specifically, if $\det(\Lambda_n^p) \geq 2\det(\Lambda^p)$ then by Corollary 2, $\det(\Lambda_n) \geq \det(\Lambda_n^p)/1.11 \geq 2\det(\Lambda^p)/1.11 \geq 2 \times 0.9\det(\Lambda)/1.11 \geq 1.5\det(\Lambda)$. Similarly, if $\det(\Lambda_n) > 2.5\det(\Lambda)$ then by Corollary 2, $\det(\Lambda_n^p) \geq 0.9\det(\Lambda_n) > 0.9 \times 2.5\det(\Lambda) \geq 0.9 \times 2.5\det(\Lambda^p)/1.1 \geq 2\det(\Lambda^p)$. □

With Lemma EC.4 we can upper bound the number of times condition $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ is active. Because $\|\phi_t\|_2 \leq 1$ almost surely, we have that $\det(\Lambda_{T-1}) \leq (T-1+\rho)^d$. On the other hand, $\det(\Lambda_0) = \rho^d$. Hence, the number of times $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ is satisfied is upper bounded by the number of times $\det(\Lambda_n) \geq 1.5\det(\Lambda)$, which is further bounded by $\lceil \log_{1.5}(T^d)\rceil = \lceil d\log_{1.5} T\rceil$, because $\rho \geq 1$. Comparing this with the definition of D_∞, we conclude that with probability $1 - O(T^{-1})$ each $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ event leads to a call of sub-routine PRIVATEMLE.

Our next lemma shows that with high probability, the constructed upper confidence bound $f(\phi(p,x_n)^\top\widehat{\theta}^p) + \gamma\sqrt{\phi(p,x_n)^\top[\Lambda^p]^{-1}\phi(p,x_n)}$ is a valid upper bound on $f(\phi(p,x_n)^\top\theta^*)$.

LEMMA EC.5. *For each time period n and price $p \in [0,1]$, define $r_n(p) = pf(\phi(p,x_n)^\top\theta^*)$, $\widehat{r}_n(p) = pf(\phi(p,x_n)^\top\widehat{\theta}^p)$ and $\overline{r}_n(p) = p\min\{1, \widehat{r}_n(p) + \gamma\sqrt{\phi(p,x_n)^\top[\Lambda^p]^{-1}\phi(p,x_n)}\}$. With probability $1 - O(T^{-1})$ the following holds for every time period t and price $p \in [0,1]$:*

$$r_n(p) \leq \overline{r}_n(p) \leq r_n(p) + 2\gamma\sqrt{\phi(p,x_n)^\top[\Lambda^p]^{-1}\phi(p,x_n)}.$$

Proof of Lemma EC.5. Because $r_n(p) \in [0,1]$ and $p \in [0,1]$ almost surely, we only need to prove that $|r_n(p) - \widehat{r}_n(p)| \leq \gamma\sqrt{\phi(p,x_n)^\top[\Lambda^p]^{-1}\phi(p,x_n)}$ for all n and p. Fix $\phi = \phi(p,x_n)$ for some $p \in [0,1]$. Decompose $|f(\phi^\top\theta^*) - f(\phi^\top\widehat{\theta}^p)|$ as

$$\left|f(\phi^\top\theta^*) - f(\phi^\top\widehat{\theta}^p)\right| \leq K\left|\phi^\top(\widehat{\theta}^p - \theta^*)\right| \leq K\|\phi\|_{[\Lambda^p]^{-1}}\|\widehat{\theta}^p - \theta^*\|_{\Lambda^p} \tag{EC.16}$$

$$\leq K(sK\sqrt{3d\ln T} + (2G+3)\sqrt{\rho} + G\nu_{\varepsilon_2',\delta_2'}\sqrt{5d\ln T}) \times \sqrt{\phi^\top[\Lambda^p]^{-1}\phi} \quad \text{(EC.17)}$$

$$\leq \gamma\sqrt{\phi^\top[\Lambda^p]^{-1}\phi}. \quad \text{(EC.18)}$$

Here, the first inequality in Eq. (EC.16) holds because $\|\theta^*\|_2 \leq 1$ and $\|\widehat{\theta}_n^p\|_2 < 2$ with high probability, and therefore $|f(\phi^\top\theta^*) - f(\phi^\top\widehat{\theta}^p)| = |f'(\phi^\top\widetilde{\theta})\phi^\top(\widehat{\theta}^p - \theta^*)| \leq K|\phi^\top(\widehat{\theta}^p - \theta^*)|$ by the mean value theorem. Eq. (EC.17) holds with probability $1 - O(T^{-1})$ by invoking Lemma 2. This completes the proof of Lemma EC.5. □

To prove the regret upper bound in Theorem 1 we need two additional technical lemmas. Both lemmas are proved in the seminal work of Abbasi-Yadkori et al. (2011) and we omit their proofs in this paper.

Lemma EC.6 **(Lemma 12, (Abbasi-Yadkori et al. 2011))**. *Let A, B be positive semi-definite matrices. Then $\sup_{\phi\neq 0}\frac{\phi^\top(A+B)\phi}{\phi^\top A\phi} \leq \frac{\det(A+B)}{\det A}$.*

Lemma EC.7 **(Lemma 11, (Abbasi-Yadkori et al. 2011))**. *For $\rho \geq 1$ and $\|\phi_n\|_2 \leq 1$ for all n, it holds that $\sum_{n=1}^T \phi_n^\top\Lambda_n^{-1}\phi_n \leq 2\ln\frac{\det(\Lambda_T)}{\det(\rho I)} \leq 2d\ln T$, where $\Lambda_n = \sum_{t<n}\phi_t\phi_t^\top + \rho I$.*

We are now ready to prove Theorem 1.

Proof of Theorem 1. With Lemma EC.5 and the definition of offered price p_n in Algorithm 1, we have with probability $1 - O(T^{-1})$ that

$$\sum_{n=1}^T p_n^* f(\phi(p_n^*, x_n)^\top\theta^*) - p_n(f(\phi(p_n, x_n)^\top\theta^*)) = \sum_{n=1}^T r_n(p_n^*) - r_n(p_n)$$

$$\leq \sum_{n=1}^T \overline{r}_n(p_n^*) - \overline{r}_n(p_n) + \overline{r}_n(p_n) - r_n(p_n) \quad \text{(EC.19)}$$

$$\leq \sum_{n=1}^T \overline{r}_n(p_n) - r_n(p_n), \quad \text{(EC.20)}$$

where Eq. (EC.19) holds because $r_n(\cdot) \leq \overline{r}_n(\cdot)$ with high probability, and Eq. (EC.20) holds because p_n is the maximizer of $\overline{r}_n(\cdot)$. Invoking Lemma EC.5 we have with probability $1 - O(T^{-1})$ that

$$\overline{r}_n(p_n) - r_n(p_n) \leq 2\gamma\sqrt{\phi_n^\top[\Lambda^p]^{-1}\phi_n}, \quad \forall n, \quad \text{(EC.21)}$$

where $\phi_n = \phi(x_n, p_n)$.

Recall that Λ^p is the differentially private sample covariance copy kept by Algorithm 1, which may or may not be updated at time n. $\Lambda_n = \sum_{t<n}\phi_t\phi_t^\top + \rho I$, on the other hand, is the true sample

covariance at time n. The algorithm pseudo-code and Lemma EC.4 shows that, with probability $1-O(T^{-1})$, one must have $\det(\Lambda^p)\geq 0.5\det(\Lambda_n^p)\geq 0.45\det(\Lambda_n)$. Subsequently, by Lemma EC.6,

$$\phi_n^\top[\Lambda^p]^{-1}\phi_n \leq \frac{1}{0.45}\phi_n^\top[\Lambda_n]^{-1}\phi_n \leq 2.3\phi_n^\top\Lambda_n^{-1}\phi_n. \tag{EC.22}$$

Combining Eqs. (EC.20,EC.21,EC.22) and invoking Lemma EC.7, we obtain

$$\begin{aligned}\sum_{n=1}^T r_n(p_n^*)-r_n(p_n) &\leq \sum_{n=1}^T 2\gamma\sqrt{2.3\phi_n^\top\Lambda_n^{-1}\phi_n} \leq 2\gamma\sqrt{2.3T}\times\sqrt{\sum_{n=1}^T\phi_n^\top\Lambda_n^{-1}\phi_n}\\ &\leq 2\gamma\sqrt{4.6dT\ln T}.\end{aligned}$$

Plugging in the scalings of γ and ρ we complete the proof of Theorem 1. □

EC.6. Proof of Theorem 2

The proof of Theorem 2 is largely the same with the proof of Theorem 1, except for the application of the second upper bound in Lemma 2. The following lemma establishes (with high probability) a lower bound on the smallest eigenvalue of $\Sigma_{T_0+1}=\sum_{t\leq T_0}\phi_t\phi_t^\top$, which is a condition for the second upper bound in Lemma 2.

Lemma EC.8. *Suppose Assumption 1 holds and T_0 is set as instructed in Theorem 2. Then for any $T\geq e^{\kappa_x^{-2}}$, with probability $1-O(T^{-1})$ it holds that $\lambda_{\min}(\Sigma_{T_0+1})\geq[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2$.*

Proof of Lemma EC.8. Because $\{p_t\}$ are chosen uniformly at random from $[0,1]$, it holds that $\mathbb{E}[\Sigma_{T_0+1}]=\sum_{t\leq T_0}\mathbb{E}[\phi_t\phi_t^\top]\succeq\kappa_x T_0 I_d$. On the other hand, because $\|\phi_t\|_2\leq 1$ almost surely, by matrix Hoeffding's inequality (see, e.g., Theorem 1.3 of (Tropp 2012)) we have that with probability $1-O(T^{-1})$,

$$\left\|\Sigma_{T_0+1}-\mathbb{E}[\Sigma_{T_0+1}]\right\|_{\mathrm{op}}\leq\sqrt{8T_0\ln(dT)}.$$

With $T_0\geq 32\kappa_x^{-2}\ln(dT)$, it holds with probability $1-O(T^{-1})$ that $\|\Sigma_{T_0+1}-\mathbb{E}[\Sigma_{T_0+1}]\|_{\mathrm{op}}\leq\frac{\kappa_x}{2}T_0$ and subsequently $\lambda_{\min}(\Sigma_{T_0+1})\geq\frac{\kappa_x}{2}T_0$. Furthermore, with $T_0\geq 2\kappa_x^{-1}[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2$, it further holds that $\frac{\kappa_x}{2}T_0\geq[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2$, satisfying the condition in Lemma 2.

The above analysis imposes the lower bound $T_0\geq\max\{32\kappa_x^{-2}\ln(dT),2\kappa_x^{-1}[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2\}$ on T_0. It is easy to verify that, with $T\geq e^{\kappa_x^{-2}}$, the condition $T_0\geq 32[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2\ln^2(dT)$ implies $T_0\geq\max\{32\kappa_x^{-2}\ln(dT),2\kappa_x^{-1}[\frac{(2G+3)\rho}{\sqrt{5d\ln T}}+\nu_{\varepsilon,\delta}G]^2\}$. This completes the proof of Lemma EC.8. □

With Lemma EC.8, Lemma 2 shows that Lemma EC.5 holds by replacing γ with $\gamma=4sK^2\sqrt{d\ln T}$. The rest of the proof of Theorem 1 remains unchanged for the proof of Theorem 2.