Privacy-Preserving Contextual Personalized Dynamic Pricing for Ride-Hailing Platforms

Bing Song[a], *Sisi Jian*[a,*]
[a] *Department of Civil and Environmental Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong SAR*
[*] *Corresponding author, Email: cesjian@ust.hk*

February 27, 2025

Keywords: Mobility service, Contextual personalized dynamic pricing, Differential privacy, Privacy-preserving, Online convex optimization

1 INTRODUCTION

In the rapidly evolving ride-hailing industry, pricing strategy is crucial in shaping market dynamics and consumer behavior. Effective pricing influences platform profitability, user perception, and demand elasticity. Dynamic pricing, which adjusts fares in real-time based on demand fluctuations, time of day, and external factors, is increasingly adopted. By leveraging advanced algorithms and big data analytics, ride-hailing platforms can predict demand patterns and optimize pricing. This adaptability maximizes revenue during peak periods and enhances customer experience with competitive off-peak fares.

Despite the advantages of dynamic pricing, designing an optimal online pricing algorithm for ride-hailing platforms is a significant challenge (Tang *et al.*, 2020). The market's complexity requires balancing multiple objectives, including revenue maximization, customer satisfaction, and driver availability. Accurately predicting demand is a major obstacle, as it involves historical ride patterns, real-time traffic conditions, and external factors like emergencies. Even minor inaccuracies can lead to lost revenue or customer dissatisfaction. Additionally, the algorithm must quickly adjust prices in response to sudden market changes, such as demand surges from special events or unexpected drops in driver supply.

Privacy risks present another critical challenge (Lei *et al.*, 2023). Competing platforms may impersonate passengers to infer the pricing strategies of ride-hailing services, potentially engaging in predatory pricing that leads to economic losses. Additionally, the leakage of passenger information can have even more severe consequences; third-party agents can deduce sensitive personal information by observing price fluctuations. For example, if a price increase occurs after a passenger completes a ride while driver supply is adequate, the agent may conclude that the passenger exhibits low price sensitivity. Such insights can be exploited for targeted advertising, discriminatory pricing, or even illicit activities, resulting in financial and privacy-related harm for the passenger.

To address these challenges, we extend the framework of Chen *et al.* (2019) by approximating personalized dynamic ride-hailing pricing as a full-information online learning problem. The binary feedback (accept/reject service) leads to non-convex loss functions. We propose an alternative formulation and connect it to contextual dynamic pricing. We then introduce techniques from differential privacy (DP) online learning and the tree-based aggregation protocol (TBAP) (Chan *et al.*, 2011) to solve this problem while ensuring user privacy. Theoretically, we demonstrate that the proposed algorithm satisfies ϵ-DP and achieves an expected regret of $\widetilde{O}(\frac{\sqrt{dT}}{\epsilon})$.

2 Problem Statement

Consider the pricing problem faced by a ride-hailing platform that matches drivers with passengers based on varying orders and collects a commission for its services. During each time step $t = 1, 2, \ldots, T$, a passenger arrives with the intention of traveling to a designated destination. We model the arrival of passengers as a stochastic process drawn from an unknown distribution, where the sequence of contexts may be adversarial in nature. Each request is characterized by a feature vector $\mathbf{x}_t \in \mathbb{R}^d$. which may include publicly observable attributes (e.g., distance and time) as well as information known only to the current passenger and the platform (e.g., the passenger's wait time). The passenger's valuation of the ride-hailing service depends on both service features and their private preferences. Assuming that the value a passenger assigns to the service is a linear function of the service features $\mathbf{x}$ and their individual preferences $\boldsymbol{\theta}$, we can express the valuation function as follows:

$$f_t(\mathbf{x}_t) = \langle \boldsymbol{\theta}, \mathbf{x}_t \rangle + s_t \tag{1}$$

In this scenario, $\boldsymbol{\theta}$ signifies a characteristic that is relevant to all passengers but remains unknown to the platform. Additionally, s_t is characterized as a scalar random variable, known as preference shocks, which correspond to the individual passenger at time t. These preference shocks are presumed to be independently and identically distributed (i.i.d.) with an average value of zero across $\mathbb{R}$. We denote the cumulative distribution function as H and the probability density function as $h(z) = H'(z)$. At each time period t, the platform sets a price p_t for the passenger. If $p_t \leq f_t(\mathbf{x}_t)$, the passenger accepts the service, generating revenue p_t for the platform. Conversely, if $p_t \geq f_t(\mathbf{x}_t)$, the passenger declines the service, resulting in no revenue. The platform's objective is to develop a pricing strategy that maximizes revenue. Crucially, the platform is equipped with previous feedback (service accepted or not accepted) at each stage and can leverage this information to adaptively modify the current price.

The objective of this study is to propose a dynamic personalized pricing strategy that not only maximizes the expected revenue of the ride-hailing platform but also ensures that the platform's pricing strategy and each passenger's valuation information remain confidential. To achieve this, we introduce the concept of DP to ensure privacy protection while providing a definition of regret for the platform to measure its revenue.

Definition 1. *(ϵ-differential privacy (Dwork et al., 2006)). A pricing policy $\mathcal{B}$ of the ride-hailing platform outputs a sequence of prices $\mathbf{P} = \{p_2, ..., p_T\} \in \mathbb{R}^{T-1}$ based on a sequence of passengers' preference shocks $\mathbf{S} = \{s_1, ..., s_T\}$ and an random sequence of service features $\mathbf{X} = \{\mathbf{x}_1, ..., \mathbf{x}_T\}$. If for any two neighboring databases $(\mathbf{X}, \mathbf{S})^a$ and $(\mathbf{X}, \mathbf{S})^b$ that differ in at most one entry in $\mathbf{S}$, and for all $\mathcal{P} \subseteq \mathbb{R}^{T-1}$, a randomized pricing policy $\mathcal{B}$ holds:*

$$Pr(\mathcal{B}((\mathbf{X}, \mathbf{S})^a)) \in \mathcal{P}) \leq e^{\epsilon} Pr(\mathcal{B}((\mathbf{X}, \mathbf{S})^b)) \in \mathcal{P}) \tag{2}$$

then this pricing policy $\mathcal{B}$ is ϵ-DP.

Platforms' Regret. The platform's utility is evaluated through the concept of regret, which quantifies the maximum expected revenue loss relative to an optimal strategy that can retrospectively identify the hidden model parameter $\boldsymbol{\theta}$. It is important to highlight that the anticipated revenue generated from a posted price (p) can be expressed as:$p \cdot \Pr(p \leq f_t) = p \cdot (1 - H(p - \boldsymbol{\theta} \cdot \mathbf{x}_t))$

By applying the first-order condition, we can derive the optimal price $p^*(\mathbf{x}_t)$:

$$p^*(\mathbf{x}_t) = \frac{1 - H(p^*(\mathbf{x}_t) - \langle \mathbf{x}_t, \boldsymbol{\theta} \rangle)}{h(p^*(\mathbf{x}_t) - \langle \mathbf{x}_t, \boldsymbol{\theta} \rangle)} \tag{3}$$

Define the virtual valuation function of the passenger $\Omega(f) = f - \frac{1-H(f)}{h(f)}$ and define the optimal pricing function $\Lambda(f) = f + \Omega^{-1}(-f)$. Then Ω is injective and Λ is well-defined and non-negative .

We can define the optimal price and maximum regret of $\mathcal{A}$ as follows:

$$p_t^* = \Lambda(\langle \mathbf{x}_t, \boldsymbol{\theta} \rangle) \tag{4}$$

$$Regret_{\mathcal{A}}(T) = \sup_{\mathbf{X}, \boldsymbol{\theta}} \sum_{t=1}^{T} (p_t^* \mathbb{I}(f_t \geq p_t) - p_t \mathbb{I}(f_t \geq p_t)) \tag{5}$$

3 Our algorithm

As we deifne the regret in Eq.(5), The loss function l_t is defined as the negative of the revenue generated during time step t, expressed as $l_t = \mathbb{I}(p_t \leq o_t)p_t$. Since this is a non-convex function and by obtaining an accurate estimate of the hidden parameter $\boldsymbol{\theta}$, we can calculate the optimal price using Equation (3). Therefore, we can reformulate the problem as an alternative online convex optimization problem using the negative log-likelihood function in Equation (6).

$$l_t(\boldsymbol{\theta}) = -log(1 - H(p - \boldsymbol{\theta} \cdot \mathbf{x}_t))\mathbb{I}\{c_t = 1\} - log(H(p - \boldsymbol{\theta} \cdot \mathbf{x}_t))\mathbb{I}\{c_t = 0\} \tag{6}$$

where $c_t \in \{0, 1\}$ means whether the passenger accepts the service after receiving the price p_t offered by the platform at time t.

We present the steps to solve Eq.(6) in Algorithm 1, followed by a detailed explanation.

Algorithm 1 Privacy-Preserving Contextual Personalized Pricing

Inputs:
1. Trip features $\{\mathbf{x_t}\}_{t\geq 1}$ 2. Set $\boldsymbol{\Theta}$, pricing function $\Lambda(\cdot)$ 3. privacy budget ϵ, Lipschitz parameter u_F, strong convexity parameter K
Algorithm Steps:
4. Set an arbitrary $\hat{\boldsymbol{\theta}}_{\mathbf{1}} \in \boldsymbol{\Theta}$ and price $p_1 = 0$
5. $\hat{\boldsymbol{\tau}}_1 \leftarrow TBAP(u_F, \epsilon, \nabla l_1^K(\hat{\theta}_1))$
for t = 1, ... **do**
    6. $\hat{\theta}_{t+1} \leftarrow argmin_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \langle \hat{\tau}_t, \boldsymbol{\theta} \rangle + \frac{K}{2} \sum_{\omega=1}^{t} ||\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\omega||^2$
    7. $p_{t+1} \leftarrow \Lambda(\langle x_{t+1}, \hat{\boldsymbol{\theta}}_{t+1} \rangle)$, observe y_{t+1}
    8. $\tau_{t+1} \leftarrow TBAP(u_F, \epsilon, \nabla l_{t+1}^K(\hat{\theta}_{t+1}))$
end for

We start to guarantee strong of Eq.(6) by imposing regularization. The K-regularized function is defined in Eq.(7).

$$l_t^K(\boldsymbol{\theta}) = l_t(\boldsymbol{\theta}) + \frac{K}{2}||\boldsymbol{\theta}||^2 \tag{7}$$

where K represents a parameter that we will modify to optimize the algotithm performance. We can ensure K-strongly convexity of each l_t^K.

The pivotal step involves utilizing the quadratic approximations $\widetilde{l}_1^K, ..., \widetilde{l}_T^K$ of the loss functions $l_1^K, ..., l_T^K$. To calculate the price p_t according to Eq.(3), this algorithm will provide an estimate $\hat{\boldsymbol{\theta}}$ of the parameter $\boldsymbol{\theta}$ at each time step t. Since l_t^K is strongly convex, the lower bound of l_t^K can be calculated by Eq.(8). Clearly, the value and gradient of l_t^K and $\widetilde{l}_t^K$ are the same at $\hat{\boldsymbol{\theta}}_t$ (that is $l_t(\hat{\boldsymbol{\theta}}_t) = \widetilde{l}_t(\hat{\boldsymbol{\theta}}_t)$ and $\nabla l_t(\hat{\boldsymbol{\theta}}_t) = \nabla \widetilde{l}_t(\hat{\boldsymbol{\theta}}_t)$).

$$\widetilde{l}_t^K = l_t^K(\hat{\boldsymbol{\theta}}_t) + \langle \nabla l_t^K(\hat{\boldsymbol{\theta}}_t), \boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t \rangle + \frac{K}{2}||\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t||^2 \tag{8}$$

Let $\widetilde{\boldsymbol{\theta}}_{t+1} = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{argmin} \sum_{\omega=1}^{t} \widetilde{l}_\omega^K(\boldsymbol{\theta})$ represent the "leader" associated with the cost functions $\widetilde{l}_1^K, ..., \widetilde{l}_t^K$. The minimization of the sum of $\widetilde{l}_t(\boldsymbol{\theta})$ is equivalent to minimizing the sum of $\widetilde{l}_t^K(\boldsymbol{\theta}) -$

$l_t^K(\hat{\boldsymbol{\theta}}_t)$, as altering a constant term does not affect the minimizer. Therefore, we can express $\widetilde{\boldsymbol{\theta}}_{t+1}$ in Eq.(9).

$$\widetilde{\boldsymbol{\theta}}_{t+1} = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{argmin} \langle \sum_{\omega=1}^{t} \nabla l_\omega^K(\hat{\boldsymbol{\theta}}_\omega), \boldsymbol{\theta} \rangle + \frac{K}{2} \sum_{\omega=1}^{t} ||\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\omega||^2 \tag{9}$$

As long as the estimate $\hat{\boldsymbol{\theta}}_t$ solely determined by the cumulative gradients $\tau_t = \sum_{\omega=1}^{t} \nabla l_\omega^K(\hat{\boldsymbol{\theta}}_\omega)$ satisfies DP, the algorithm output p_t, which pertains to passenger preference shocks s_t, maintains DP. We utilize TBAP to accurately compute the differentially private version $\hat{\tau}_t$ of the cumulative gradient information τ_t. Substituting $\hat{\tau}_t$ into Eq.(9) yields Eq.(10).

$$\hat{\boldsymbol{\theta}}_{t+1} = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{argmin} \langle \sum_{\omega=1}^{t} \hat{\boldsymbol{\tau}}_t, \boldsymbol{\theta} \rangle + \frac{K}{2} \sum_{\omega=1}^{t} ||\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_\omega||^2 \tag{10}$$

Simply by substituting the result of Eq.(10), the current parameter estimate $\hat{\boldsymbol{\theta}}_{t+1}$, into Eq.(3), the platform can compute the price p_{t+1} that the platform should provide to the passengers at time $t+1$.

Theorem 1. ***Privacy Guarantee.*** *Algorithm 1 is ϵ-DP for any trip features and passenger preference shock sequences $(\mathbf{x}_1, s_1), ..., (\mathbf{x}_T, s_T)$.*

Theorem 2. *Given the valuation function described in Equation (1), Algorithm 1 attains regret on the scale of $\widetilde{O}(\sqrt{dT}/\epsilon)$.*

4 Conclusion

This study proposes a novel differentially private contextual dynamic pricing method to address privacy concerns in the dynamic pricing landscape of ride-hailing platforms. By integrating DP—a gold standard in privacy preservation—with online convex optimization, the approach effectively introduces noise injection to protect passengers' sensitive information. We consider a scenario where passengers receive prices from the platform and decide whether to accept the service, assuming a linear valuation function. We transform this non-convex setting into an online convex problem, then leverage the concept of DP and TBAP algorithms to propose our method, balancing privacy protection and data utility through noise injection in the cumulative gradients. The proposed approach meets the stringent requirements of differential privacy, ensuring robust protection of passenger information while establishing an upper limit on regret. The expected regret is approximately $(\widetilde{O}(\frac{\sqrt{dT}}{\epsilon}))$, indicating that as the potential passenger pool (T) increases, the algorithm achieves optimal performance for individual passengers and nearly cost-free reinforcement of their privacy.

References

Chan, T-H Hubert, Shi, Elaine, & Song, Dawn. 2011. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)*, **14**(3), 1–24.

Chen, Haipeng, Jiao, Yan, Qin, Zhiwei, Tang, Xiaocheng, Li, Hao, An, Bo, Zhu, Hongtu, & Ye, Jieping. 2019. InBEDE: Integrating contextual bandit with TD learning for joint pricing and dispatch of ride-hailing platforms. *Pages 61–70 of: 2019 IEEE International Conference on Data Mining (ICDM)*. IEEE.

Dwork, Cynthia, McSherry, Frank, Nissim, Kobbi, & Smith, Adam. 2006. Calibrating noise to sensitivity in private data analysis. *Pages 265–284 of: Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*. Springer.

Lei, Yanzhe, Miao, Sentao, & Momot, Ruslan. 2023. Privacy-preserving personalized revenue management. *Management Science*.

Tang, Wei, Ho, Chien-Ju, & Liu, Yang. 2020. Differentially Private Contextual Dynamic Pricing. *Pages 1368–1376 of: AAMAS*.