



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Privacy-Preserving Dynamic Personalized Pricing with Demand Learning

Xi Chen, David Simchi-Levi, Yining Wang

To cite this article:

Xi Chen, David Simchi-Levi, Yining Wang (2021) Privacy-Preserving Dynamic Personalized Pricing with Demand Learning. Management Science

Published online in Articles in Advance 11 Oct 2021

. <https://doi.org/10.1287/mnsc.2021.4129>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2021, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Privacy-Preserving Dynamic Personalized Pricing with Demand Learning

Xi Chen,^{a,*} David Simchi-Levi,^b Yining Wang^c

^aLeonard N. Stern School of Business, New York University, New York, New York 10012; ^bMIT Institute for Data, Systems, and Society, Department of Civil and Environmental Engineering, and Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^cWarrington College of Business, University of Florida, Gainesville, Florida 32611

*Corresponding author

Contact: xc13@stern.nyu.edu,  <https://orcid.org/0000-0002-9049-9452> (XC); dslevi@mit.edu (DS-L); yining.wang@warrington.ufl.edu (YW)

Received: October 4, 2020

Revised: December 27, 2020; March 28, 2021; May 16, 2021

Accepted: May 18, 2021

Published Online in Articles in Advance: October 11, 2021

<https://doi.org/10.1287/mnsc.2021.4129>

Copyright: © 2021 INFORMS

Abstract. The prevalence of e-commerce has made customers' detailed personal information readily accessible to retailers, and this information has been widely used in pricing decisions. When using personalized information, the question of how to protect the privacy of such information becomes a critical issue in practice. In this paper, we consider a dynamic pricing problem over T time periods with an *unknown* demand function of posted price and personalized information. At each time t , the retailer observes an arriving customer's personal information and offers a price. The customer then makes the purchase decision, which will be utilized by the retailer to learn the underlying demand function. There is potentially a serious privacy concern during this process: a third-party agent might infer the personalized information and purchase decisions from price changes in the pricing system. Using the fundamental framework of differential privacy from computer science, we develop a privacy-preserving dynamic pricing policy, which tries to maximize the retailer revenue while avoiding information leakage of individual customer's information and purchasing decisions. To this end, we first introduce a notion of *anticipating* (ϵ, δ) -differential privacy that is tailored to the dynamic pricing problem. Our policy achieves both the privacy guarantee and the performance guarantee in terms of regret. Roughly speaking, for d -dimensional personalized information, our algorithm achieves the expected regret at the order of $\tilde{O}(\epsilon^{-1}\sqrt{d^3T})$ when the customers' information is adversarially chosen. For stochastic personalized information, the regret bound can be further improved to $\tilde{O}(\sqrt{d^2T} + \epsilon^{-2}d^2)$.

History: Accepted by J. George Shanthikumar, big data analytics.

Funding: X. Chen is supported by the National Science Foundation [Grant IIS-1845444]. D. Simchi-Levi is supported by MIT-Accenture Alliance for Business Analytics and the MIT Data Science Laboratory.

Supplemental Material: The supplementary material is available at <https://doi.org/10.1287/mnsc.2021.4129>.

Keywords: differential privacy (DP) • dynamic pricing • generalized linear bandits • personal information

1. Introduction

The increasing prominence of e-commerce has given retailers an unprecedented power to understand customers as individuals and to tailor their services accordingly. For example, personal information is known to be used in pricing on travel websites (Hannak et al. 2014) and Amazon (Chen et al. 2016); Linden et al. (2003) illustrate how personal information is used in Amazon recommender systems to achieve a dramatic increase in click-through and conversion rates. Although personalized pricing may involve complicated legal issues in many domains, it has been adopted or considered in several key industries, such as air travel, hotel booking, insurance, and ride sharing. For example, according to Tringale (2018), "Hotel websites such as Orbitz (whose parent company is

Expedia) and auto dealers like Tesla utilize personalized pricing to their advantage when conducting sales with a customer. Even Uber has dabbled in personalized pricing by offering 'premium pricing' to predict which users are willing to pay more to go to a certain location." As reported by Mohammed (2017), when using Orbitz, for identical flights, hotel, and type of room, the price of the traveling package found on a laptop was 6.5% more than the price offered on the Orbitz app. Moreover, in practice, instead of directly charging different prices, the e-commerce platforms usually use the discount or promotions to implement personalized pricing strategies.

Although the availability of personal data (e.g., location, web search histories, media consumption, social media activities) enables targeted services for an

individual customer, it poses significant privacy issues in practice (e.g., Apple Differential Privacy Team 2017). Many existing privacy-protection approaches are rather ad hoc by “anonymizing” personal information. However, such ad hoc anonymization leads to two issues. First, it is difficult to quantify the level of privacy. Second, it has been shown that a deanonymization procedure can easily jeopardize privacy. Examples include the deanonymization of released AOL search logs (Barbaro and Zeller 2006) and movie-watching records in Netflix challenges (Narayanan and Shmatikov 2008). Therefore, personalized operations management urgently calls for mathematically rigorous privacy-preserving methods to prevent personal information leakage in online decision making. On one hand, personalized revenue management has received a significant amount of attention in recent operations literature (see, e.g., Cheung and Simchi-Levi (2017), Ban and Keskin (2021), and references therein). On the other hand, the question of how to protect an individual’s privacy has not been well explored in the existing literature.

In this paper, we study how to systematically protect an individual’s privacy in the dynamic pricing problem with demand learning. Given T time periods, a potential customer arrives at each time t , and the retailer receives x_t containing information about the incoming customer, such as age, location, purchase history and ratings, and credit scores, among other things. We consider a very general personalized setting where the customers are *heterogeneous*, and thus the feature $\{x_t\}_{t=1}^T$ does not necessarily follow the same distribution. By observing the personal information x_t , the retailer offers the customer a price $p_t \in [0, 1]$. The customer then makes $y_t \in \mathbb{R}$, where the random demand y_t follows a *generalized linear model* of a feature vector $\phi(x_t, p_t) \in \mathbb{R}^d$ (see (1)) and the retailer collects revenue $p_t y_t$. The objective of the retailer is to maximize the expected revenue over the entire T time periods or, more specifically, $\mathbb{E}[\sum_{t=1}^T p_t y_t]$. As this paper focuses on how to protect an individual’s sensitive information, we consider a stylized setting of pricing a *single product* with *unlimited* inventories available.

Because of the personalized nature, the aforementioned pricing procedure involves the use of individuals’ sensitive information, such as customers’ personal information, characterized by x_t and their purchase history, designated by y_t (e.g., whether a purchase was made at time t). Thanks to secured internet communication channels, the information (x_t, p_t, y_t) at time t is usually securely transmitted and thus only revealed to the retailer and the particular customer coming at time t . However, although the information at time t is not directly accessible to future customers, the sensitive information is not completely shielded from outside third-party agents (a.k.a. attackers or

adversaries) because of the ripple effects of historical customers’ data on future pricing decisions in a data-driven pricing system. Indeed, a third-party agent who observes his own posted prices *in the future* can potentially infer an individual’s personal information x_t and purchase decision y_t . We provide two examples showing how the sensitive data at time t could be potentially breached and why such privacy leakage could incur serious challenges to the integrity of the underlying pricing system.

- *Leakage of purchase activity y_t* : For sensitive commodities such as medications, customers’ purchasing decisions $\{y_t\}$ must be well protected from the public, as such purchases may potentially reveal purchasers’ underlying medical conditions. Some dynamic pricing policies would increase prices facing increased sales volumes for a higher profit. Such behavior might inadvertently leak information about y_t to a third party via the fluctuation of prices. For example, a third-party agent might place orders immediately before and after a person of interest, and if the agent sees a slight spike in his or her received prices might be able to infer the purchase decision y_t of the person of interest.

- *Leakage of customers’ personal information x_t* : When making the price decision p_t for an arriving customer at time t , the retailer makes use of the customer’s personal information x_t . Some components of x_t , such as the customer’s age, credit history, and prior purchases, are highly sensitive and should be protected. Consider a natural pricing policy that is highly “local” to personal information (e.g., posting similar prices to future customers with a similar profile to customer t). A third-party agent could arrive before and after a person of interest with guesses of personal information to detect whether there are noticeable changes in the prices. Then, the agent would be able to infer to some degree about the personal information x_t of the individual of interest.

In summary, it is vital to develop systematic and mathematically rigorous policies that *provably* protect customers’ privacy. As we previously discussed, simple data anonymization lacks a theoretical foundation and can be jeopardized. On the other hand, the notion of *differential privacy* (DP), which was proposed in the computer science field (Dwork et al. 2006a, b), has laid a solid foundation for private data analysis and achieved great success in industries. The DP not only is a gold standard notion in academia but also has been widely adopted by companies such as Apple (Apple Differential Privacy Team 2017), Google (Erlingsson et al. 2014), and Microsoft (Ding et al. 2017) and government agencies such as the U.S. Census Bureau (Abowd 2018). The aim of this paper is therefore to build on the differential privacy notion to design mathematically rigorous privacy policies with provable utility (regret) guarantees for the dynamic personalized pricing problem.

1.1. Our Contributions

The major contributions of this paper can be summarized as follows.

1.1.1. Near-Optimal Regret of Provably Privacy-Aware Pricing Policies. Building on the notion of anticipating differential privacy, we propose a privacy-aware personalized pricing algorithm that enjoys rigorous regret guarantees. More specifically, in a general setting, when the personalized information of each coming customer can be adversarially chosen, our policy achieves a regret upper bound of $\tilde{O}(\varepsilon^{-1}\sqrt{d^3T})$, where ε is the parameter in DP (a smaller ε implies a stronger privacy preservation of the resulting algorithm), d is the dimension of the feature map $\phi(x_t, p_t)$, T is the time horizon, and $\tilde{O}(\cdot)$ hides logarithmic factors (see Theorem 1). The \sqrt{T} dependency on the time horizon T in this regret upper bound is optimal (Broder and Rusmevichientong 2012).

In addition to the regret upper bound for the general personalized information setting, we also study a “stochastic” setting in which the customer’s personal information $\{x_t\}$ is assumed to be stochastic and independently and identically distributed (i.i.d.) from an unknown nondegenerate distribution. We remark that this is a common assumption/setting studied in the existing literature (Qiang and Bayati 2016, Miao et al. 2019). In this setting, with some changes of hyperparameters of our proposed algorithm, an improved regret upper bound of $\tilde{O}(d\sqrt{T} + \varepsilon^{-2}d^2)$ can be proved (see Theorem 2). One attractive property of this bound is that it separates the dependency on conventional problem parameters (i.e., d and T) from the privacy-related parameter (i.e., ε). The dominating term (with $T \rightarrow \infty$) in this regret bound—namely, the $\tilde{O}(d\sqrt{T})$ term—is optimal in both d and T , as shown in Dani et al. (2008).

In both the general setting and the stochastic setting, the regret upper bounds of either $\tilde{O}(\varepsilon^{-1}\sqrt{d^3T})$ or $\tilde{O}(d\sqrt{T} + \varepsilon^{-2}d^2)$ also characterize the trade-offs between customers’ privacy protection and the revenue (surplus) of the seller under the designed policy. More specifically, the $\varepsilon > 0$ parameter characterizes the level of customers’ privacy protection, with a smaller ε corresponding to stronger protection against malicious agents. Clearly, as both regret upper bounds depend inversely on ε , it shows that as the seller seeks stronger protection over the privacy of customers’ personalized data, the more he or she will suffer from decreased revenue (and a larger regret). This revenue loss is due to additional efforts/randomization required for data privacy protection.

Finally, the privacy requirements imposed on the seller’s policy also have interesting implications on consumer surplus. In Section 9.2 of this paper, we provide numerical results to characterize the trade-offs

between consumers’ privacy protection and consumer surplus. We find that as the implied privacy protection becomes weaker (i.e., the seller having less ability to discriminate against customers based on their personal data and features, resembling a transition from the first-degree to the third-degree price discrimination), the consumer surplus increases because the seller extracts less of the consumer surplus from his or her limited ability to carry out price discrimination.

1.1.2. Technical Contributions. Our proposed framework for privacy-preserving personalized dynamic pricing makes use of several existing privacy-aware learning/releasing techniques, such as the Analyze-Gauss method in online principal component analysis (Dwork et al. 2014), the tree-based aggregation technique for releasing serial data (Chan et al. 2011), and differentially private empirical risk minimization methods (Chaudhuri et al. 2011, Kifer et al. 2012). On the other hand, the development and analysis of our proposed method make several key technical contributions to the general topic of privacy-aware sequential decision making in revenue management problems, which we briefly summarize as follows:

1. One salient feature of this paper is the inclusion of customers’ personal information x_t as sensitive data that need to be protected, which is different from existing works (Tang et al. 2020), where only purchase activities y_t are regarded as sensitive data (see Section 2 for more discussions). The objective of protecting privacy in $\{x_t\}$ presents two technical challenges. First, as $\{x_t\}$ and, subsequently, the feature representations $\{\phi_t\}$ are sensitive data, one cannot directly apply the private follow-the-regularized-leader approach in Tang et al. (2020) to the dynamic pricing problem. Furthermore, the sensitivity of $\{x_t\}$ implies the sensitivity of $\{p_t\}$ as well, because prices offered to incoming customers must be strongly associated with customers’ personal information to achieve good revenue performances. To address these challenges, we build our DP setting on the notion of anticipating DP (Shariff and Sheffet 2018), which excludes prices in prior selling periods from the outcome sets of a randomized algorithm.

2. The demand rate function f as a function of price p and personal information x is modeled in this paper as a generalized linear model within the exponential family. Despite its apparent similarity to linear models, such generalization results in significant challenges when privacy concerns are considered. In fact, this is still an open problem for generalized linear contextual bandit under the DP guarantee. More specifically, the results of Shariff and Sheffet (2018) on privacy-aware linear bandits rely heavily on the fact that the ordinary least squares solution is in a closed form with two simple sufficient statistics: the sample covariance matrix $X^T X$ and the response-weighted feature vector $X^T y$. With

the postprocessing property of DP (which we briefly discuss in Section 4.4), it suffices to obtain privacy-preserved copies of $X^T X$ and $X^T y$ at each time. By contrast, parameter estimates in generalized linear models are usually obtained using maximum likelihood estimates (MLEs), which do not have simple sufficient statistics. It is nearly impossible to guarantee the privacy and a nontrivial regret simultaneously if the MLE is updated at every period. To overcome this challenge, we make the important observation that the required number of updates of MLEs can be reduced significantly (i.e., only $O(d \log T)$ periods of updates will be sufficient). This key observation allows us to compose differentially private empirical risk minimizers (Kifer et al. 2012) to arrive at a privacy-aware contextual bandit algorithm even without explicit sufficient statistics.

3. The generalized linear model for demand rate modeling resembles existing works on parametric contextual bandits without privacy constraints (Filippi et al. 2010, Li et al. 2017, Wang et al. 2019). One significant limitation of these existing works is that, without assuming stochasticity of the contextual vectors, the optimization of parameter estimates in these works is usually nonconvex. Examples include the robustified Z-estimation in Filippi et al. (2010) and the constrained least-squares formulation in Wang et al. (2019), both of which are nonconvex for some popular generalized linear models such as the logistic regression model. Although such nonconvexity poses only computational difficulties in nonprivate bandit algorithms, these challenges become much more significant when privacy constraints are imposed, because most existing techniques of DP stochastic optimization require convexity (Chaudhuri et al. 2011, Kifer et al. 2012), and the general privacy-aware nonconvex optimization is extremely difficult.

To overcome this challenge, this paper analyzes a constrained MLE in a more refined style with a relatively large regularization parameter, demonstrating with high probability that the solution to the constrained MLE lies in the strict interior of the constraint set (see Lemma EC.1 in the supplementary material). This result then implies the first-order Karush-Kuhn-Tucker (KKT) condition of the solution, from which the Z-estimation analysis in Li et al. (2017) and Filippi et al. (2010) can be used together with the analysis of an objective-perturbed convex minimization problem to obtain satisfactory regret upper bounds.

1.2. Organization

The rest of the paper is organized as follows. Section 2 discusses the related literature in both dynamic pricing and differential privacy. We set up our pricing models and formalize the anticipating DP in Sections 3 and 4. Our policy is presented in Section 5, which contains two components: *privacy releasers* and *price*

optimizers. Sections 6 and 7 establish the privacy and regret guarantees. Section 8 provides the numerical studies. In Section 9, we further provide some discussions on the naive input perturbation approach and insights of the impact on consumer surplus, followed by a conclusion in Section 10.

2. Literature Review

This section briefly reviews related research from both the personalized pricing and differential privacy literatures.

2.1. Personalized Dynamic Pricing with Demand Learning

Because of the increasing popularity of online retailing, dynamic pricing with demand learning has become an active research area in revenue management in the past 10 years (see, e.g., Araman and Caldentey (2009), Besbes and Zeevi (2009, 2015), Farias and Van Roy (2010), Harrison et al. (2012), Broder and Rusmevichientong (2012), den Boer and Zwart (2013), Wang et al. (2014, 2020), Chen et al. (2015), Cheung et al. (2017), and Ferreira et al. (2018)). More recently, because of the availability of abundant personal information, personalized pricing with feature information has been investigated in several works. For example, Chen et al. (2021) studied offline personalized pricing and quantified the statistical property of the MLE. Cohen et al. (2020) considered a binary thresholding model for purchasing decisions by comparing a linear function of the feature and the posted price, proposed an ellipsoid-based method for dynamic pricing, and established the worst case regret bound. Qiang and Bayati (2016) considered a linear demand model and studied the performance of the greedy iterated least squares. Ban and Keskin (2021) and Javanmard and Nazerzadeh (2019) studied the personalized dynamic pricing problem in high-dimensional settings with sparsity assumption of features. A very recent work by Tang et al. (2020) studied differentially private contextual dynamic pricing and proposed a follow-the-approximate-leader-type policy. Our work differs from this paper in several respects. First, we protect the personal information $\{x_t\}$, whereas Tang et al. (2020) treated this information as public. Second, Tang et al. (2020) adopted the classical DP notion, whereas we consider the notion of anticipating DP. Finally, we assume that the demand follows a generalized linear model of a feature map of personal information and price, whereas Tang et al. (2020) considered a binary thresholding purchase model with a linear mapping of contextual information.

2.2. Differential Privacy for Online Learning

Since the notation of (ϵ, δ) -differential privacy was proposed by Dwork et al. (2006a, b), it has become a gold standard for privacy-preserving data analysis in

both academia and industry. Please refer to the survey by Dwork and Roth (2014) for a comprehensive introduction to DP.

Built on this classical notion, other privacy notions have also been developed in the literature, such as Gaussian DP (Dong et al. 2019), joint DP (Shariff and Sheffet 2018), local DP (Evfimievski et al. 2003, Kasiviswanathan et al. 2011), average-KL (Kullback-Leibler) DP (Wang et al. 2016), and per-instance DP (Wang 2019). Our notion of anticipating DP is motivated by the joint DP (Shariff and Sheffet 2018) designed for linear contextual bandits. Although the work of Shariff and Sheffet (2018) studied the linear contextual bandits subject to differential privacy constraints, their methods and analysis are built on the noisy perturbation of sufficient statistics (namely, the sample covariance and sample average). Thus, their method is *not* applicable to the personalized pricing question, where generalized linear demand models are widely used (see also the technical challenges summarized in the introduction).

In DP, there are several fundamental techniques, such as composition, postprocessing (see Section 4.3 and Dwork and Roth (2014)), partial-sum by tree-based aggregation (Dwork et al. 2010, Chan et al. 2011), and “objective perturbation” (Chaudhuri et al. 2011, Kifer et al. 2012). In our designed personalized dynamic pricing algorithm, we build on these important techniques to make sure that our algorithm is differentially private.

The techniques of DP have been applied to multi-armed bandit problems. For example, Mishra and Thakurta (2015) developed differentially private upper confidence bound (UCB) and Thompson sampling algorithms for classical bandits. Mishra and Thakurta (2015) and Shariff and Sheffet (2018) further studied differentially private linear contextual bandits, where Mishra and Thakurta (2015) protected the privacy of rewards and Shariff and Sheffet (2018) protected both rewards and contextual information. However, for linear bandits, because the maximum likelihood estimator admits a simple closed-form solution, one only needs to protect the sufficient statistics (e.g., $\sum_{t'=1}^t x_{t'} x_{t'}^\top$ and $\sum_{t'=1}^t y_{t'} x_{t'}$). On the other hand, we consider a much more general demand model following a generalized linear model. Therefore, the corresponding MLE does not admit a closed-form solution; we address this challenge by providing a new analysis of constrained MLE properties. There are other interesting private online learning frameworks developed in recent literature. For example, the private sequential learning model was proposed in Tsitsiklis et al. (2021) (for noiseless responses) and further investigated in Xu (2018) and Xu et al. (2020) (for noisy responses). In particular, Xu et al. (2020) quantified the optimal query complexity for private sequential learning against eavesdropping. Whereas the existing

privacy literature mainly focuses on protecting a data owner’s privacy, this work investigates how to protect the privacy of a learner who sequentially queries a database and receives binary responses. We note that the goal of private sequential learning is to learn a global parameter—for example, “the highest price to charge so that at least 50% of the consumers would purchase” in pricing domain (Xu et al. 2020, p. 4)—and to make sure the adversary cannot infer the final released price. By contrast, our goal is to use sequential decision making to maximize revenue while protecting individuals’ personalized information and purchasing decisions.

In the recent work of Lei et al. (2020), an *offline* personalized pricing setting is studied with differential privacy guarantees. The recent work of Zheng et al. (2020) studied the stronger local privacy notion and derived an algorithm with $\tilde{O}(T^{3/4})$ regret bound for the generalized linear model, which is worse than the regret bounds obtained in this paper.

3. Pricing Models and Assumptions

The basic setting of personalized dynamic pricing has been described in the introduction. In this section, we provide more technical details of the problem setting. At each time t with the observed personal information x_t and the posted price p_t , the (random) demand realized by customer at time t is modeled by a generalized linear model within the exponential family, taking the form of

$$\Pr[y_t = y \mid p_t, x_t, \theta^*] = \exp\{\zeta(y\phi_t^\top \theta^* - m(\phi_t^\top \theta^*)) + h(y)\}, \quad (1)$$

where $\phi_t = \phi(x_t, p_t) \in \mathbb{R}^d$ is a known feature map, $\theta^* \in \mathbb{R}^d$ is an unknown linear model, and $\zeta, m(\cdot), h(\cdot)$ are components of the distribution family. Some examples of exponential family distributions include the Gaussian distribution and the logistic model, which are given at the end of this section. It is easy to verify that $f(\phi_t^\top \theta^*) := m'(\phi_t^\top \theta^*)$ is the expectation of y_t conditioned on p_t, x_t , and θ^* . Hence, we can equivalently write Equation (1) as

$$y_t = f(\phi_t^\top \theta^*) + \xi_t, \quad (2)$$

where $\phi_t = \phi(x_t, p_t)$ and ξ_t are independent random variables satisfying $\mathbb{E}[\xi_t \mid p_t, x_t] = 0$.

We next specify the filtration process of x_t and p_t . Let $\mathcal{F}_t = \{(x_\tau, y_\tau, p_\tau)\}_{\tau=1}^t$ be the history up to time period t . In the most general setting, the features $\{x_t\}_{t=1}^T$ of the T customers are arbitrarily chosen before the pricing process starts.¹ The price p_t at each time t is subsequently chosen by the dynamic pricing policy conditioned on filtration \mathcal{F}_{t-1} and x_t . The demand y_t is then realized via $y_t = f(\phi_t^\top \theta^*) + \xi_t$, where $\phi_t = \phi(x_t, p_t)$ and $\mathbb{E}[\xi_t \mid x_t, p_t, \mathcal{F}_{t-1}] = 0$.

Throughout this paper we impose the following conditions on the distribution family, the linear model, and the feature map:

1. There exists a parameter $B_Y < \infty$ such that $|y_t| \leq B_Y$ for all time periods t in all databases D .
2. Both the feature vectors and the linear model have at most unit norm, or more specifically, $\|\phi(x, p)\|_2, \|\theta^*\|_2 \leq 1$ for all x, p .
3. The stochastic noises $\{\xi_t\}$ are centered and sub-Gaussian, meaning that $\mathbb{E}[\xi_t | x_t, p_t, \mathcal{F}_{t-1}] = 0$ and there exists $s < \infty$ such that $\mathbb{E}[e^{\lambda \xi_t} | x_t, p_t, \mathcal{F}_{t-1}] \leq e^{\lambda^2 s^2 / 2}$ for all $\lambda \in \mathbb{R}$.
4. The expectation $f(\cdot) = m'(\cdot)$ maps \mathbb{R} to $[0, 1]$ is continuously differentiable and strictly monotonically increasing. Furthermore, for all $|z| \leq 2$, $K^{-1} \leq f'(z) \leq K$ for some constant $1 \leq K < \infty$.
5. In Equation (1), ζ satisfies $G^{-1} \leq \zeta \leq G$ for some constant $1 \leq G < \infty$.

We give some common examples that fall into Equation (1) and satisfy all imposed conditions.

Example 1 (Gaussian Model).

In the Gaussian model, the realized demand y_t follows $y_t = \phi_t^\top \theta^* + \xi_t$ with $\xi_t \sim \mathcal{N}(0, 1)$. It is easy to verify that the Gaussian model falls into Equation (2) with $\zeta = 1$, $m(z) = \frac{1}{2}z^2$, $f(z) = m'(z) = z$, and $h(y) = -\frac{1}{2}y^2 - \frac{1}{2}\ln(2\pi)$. The Gaussian model also satisfies all imposed conditions with high probability with $B_Y \leq s\sqrt{\ln T}$, $s = 1$, $K = 1$, and $G = 1$.

Example 2 (Logistic Model).

In the logistic model, the realized demand y_t is supported on $\{0, 1\}$, following the logistic distribution $\Pr[y_t = 1 | \phi_t, \theta^*] = e^{\phi_t^\top \theta^*} / (1 + e^{\phi_t^\top \theta^*})$. It is easy to verify that the logistic model falls into Equation (2) with $\zeta = 1$, $m(z) = \ln(1 + e^z)$, $f(z) = m'(z) = e^z / (1 + e^z)$, and $h(y) = 1$. The logistic model also satisfies all imposed conditions with $B_Y = 1$, $s = 1$, $K = (1 + e^2)/e^2$, and $G = 1$.

4. Preliminaries on Differential Privacy

In this section we present background material on differential privacy, the core privacy concept adopted in this paper. We start with the introduction of the standard differential privacy concept, and then we show how the privacy concept could be extended to its “anticipating” version, which is more appropriate for data-driven sequential decision-making problems. Finally, we discuss two fundamental concepts of composition and postprocessing, which are essential in designing complex differentially private systems. For a full technical treatment and historical motivations, readers are referred to the comprehensive review by Dwork and Roth (2014).

4.1. Differential Privacy


Differential privacy is a mathematically rigorous measure of privacy protection and has been extensively

studied and applied since its proposal in the work of Dwork et al. (2006b). At a higher level, the fundamental concept behind differential privacy is the *impossibility* of distinguishing two “neighboring databases” (differing only on a single entry) with high probability, based on publicly available information about the database. To facilitate such probabilistic indistinguishability, the conventional approach is to artificially calibrate *stochastic noise* into the process or the outputs of differentially private algorithms.

More specifically, Figure 1 gives an intuitive illustration of the differential privacy concept applied to our dynamic personalized pricing problem. Suppose at time t the incoming customer with the context vector x_t is being offered price p_t and makes purchase decision y_t . The price decisions $\{p_t\}_{t=1}^T$ produced by the pricing algorithm are usually random, and therefore we can use P to denote the joint distribution of these prices. The concept of differential privacy requires that if a customer’s personal data change from (x_t, y_t) to (x'_t, y'_t) while all the other $T - 1$ customers’ data remain unchanged, the joint distributions of the posted prices P will change to a distribution Q that is very close to P . The closer P and Q are under the hypothetical personal data change $(x_t, y_t) \rightarrow (x'_t, y'_t)$, the better data privacy is protected under the pricing policy.

Why is the close proximity of price distributions P and Q a good measurement of a pricing algorithm’s privacy protection? Assume that a malicious agent would like to extract the sensitive information of a particular customer of interest, who arrives in the system at time t . The malicious agent must extract such sensitive data based solely on *publicly available* information, which in this case would be the firm’s posted prices p_1, \dots, p_T . Here, “public information” in the differential privacy literature refers to the information or released data that can be accessed by a malicious adversary, because these data are used by the adversary to infer the personalized data of the customers, whose privacy is to be protected. If the price distributions P and Q produced by the pricing algorithm are very

Figure 1. (Color online) Illustration of the Differential Privacy Concept



similar, then it is *information-theoretically* not possible for the malicious agent to distinguish with reasonable success probability between a customer (x_t, y_t) and another hypothetical customer (x'_t, y'_t) (see Figure 1). This means that no matter how smart the malicious agent is, it is impossible for that agent to extract very much sensitive data from the customer of interest simply based on publicly available price information.

Mathematically speaking, we use D to denote the database of all sensitive data $\{(x_t, y_t)\}_{t=1}^T$ for all of the T customers. For convenience of presentation, we also write $o_t = (x_t, y_t)$. A database D' is a *neighboring database* of D if and only if D' and D only differ at a single time period. More specifically, $D = \{o_t\}_{t=1}^T, D' = \{o'_t\}_{t=1}^T$ are neighboring databases if there exists t such that $o_t \neq o'_t$ and $o_\tau = o'_\tau$ for all $\tau \neq t$. Suppose a pricing algorithm A operates with input database D and produces randomized price output $A(D) = (p_1, \dots, p_T)$. The following definition gives a rigorous formulation of (ϵ, δ) -differential privacy.

Definition 1 ((ϵ, δ) -Differential Privacy (Dwork et al. 2006b)). For $\epsilon, \delta > 0$, a randomized algorithm A satisfies (ϵ, δ) -differential privacy if for every pair of neighboring databases D, D' and measurable set $\mathcal{A} \subseteq [\underline{p}, \bar{p}]^T$, it holds that

$$\Pr[A(D) \in \mathcal{A}] \leq e^\epsilon \Pr[A(D') \in \mathcal{A}] + \delta.$$

To facilitate the understanding of this definition, we explain why the multiplicative factor e^ϵ is critical and the role of the parameter δ in practice. Let us first explain why the DP-definition in Definition 1 adopts a multiplicative factor e^ϵ rather than an additive bound of $|\Pr[A(D) \in \mathcal{A}] - \Pr[A(D') \in \mathcal{A}]|$. Imagine that two neighboring data sets D, D' give rise to the same output O with probabilities $p_1 = \Pr[O | D]$ and $p_2 = \Pr[O | D']$. The key is to prevent a malicious party from distinguishing between D and D' based on the observation of O . If an additive guarantee is involved $|p_1 - p_2| \leq \epsilon$, then it is possible that $p_1 = 0$ and $p_2 = \epsilon$. If this is the case, the adversary would be 100% *sure whether the underlying data set is D or D'* once she observes the output O (because $p_1 = 0$ implies that it is *impossible* to observe O given D). This means that with probability ϵ , which is usually not that small (e.g., $\epsilon = 0.1$), a *catastrophe* (i.e., an outside adversary being *completely certain* about the customer's private data) will occur with 10% probability. On the other hand, if the guarantee is multiplicative (e.g., $0.9p_2 \leq p_1 \leq 1.1p_2$) then the adversary *cannot* completely distinguish between D and D' no matter how small p_1 or p_2 is. Following this discussion on the multiplicative factor versus the additive factor, because δ is an additive term, it corresponds to the probability of a catastrophe happening that allows the adversary to completely infer the privacy information about customers' data.

Because we do not want a catastrophe to happen, δ needs to be set *overwhelmingly* small. With a tiny δ value in the DP-definition, more specifically, the adversary is *always* able to conclude that D (or D') is more likely than the other, but such preference of likelihood is never going to exceed a ratio of e^ϵ . For example, with $\epsilon = 0.1$, the adversary may conclude that D is 10.5% more possible than D' based on observations of published data O but will never be able to completely/deterministically distinguish D from D' based on O .

4.2. Anticipating Differential Privacy

Despite being a widely adopted measure, the DP notion as stated in Definition 1 cannot be directly applied to dynamic pricing for several reasons. First, Definition 1 would not lead to useful pricing policies. This is because, essentially, Definition 1 requires that conditioned on the output of the *entire* posted price sequence, the adversary cannot distinguish between o_t and o'_t in a probabilistic sense. On the other hand, for high-profit personalized pricing policies, once the customer's personal information x_t changes, the price p_t offered to that customer must change accordingly in order to achieve high expected revenue, making inference of x_t much easier given p_t . Furthermore, as we have discussed in the previous paragraphs, the communications of (x_t, p_t, y_t) at time t are secured in practice, and therefore, an adversary should not have the capability of accessing the price p_t at time t . From this perspective, the classical DP notion defined in Definition 1 is too strong because it implicitly allows the adversary to access the price at time t (as p_t belongs to the output $A(D)$). In a practical setting, however, the adversary is only able to access information during other time periods (e.g., by maliciously sending fake customers to obtain price quotes) to infer the sensitive information about an individual at time t . In other words, in the following anticipating DP definition (see Definition 2), the price offered to a specific customer of interest p_t is not public information, as we can expect basic communication security between the customer and the seller. However, prices offered to *other* customers are considered public information because a malicious adversary could pretend to be a customer and extract such price information and subsequently infer the private data of the customer of interest based on such extracted price information.

This argument can be made rigorous by the following proposition. The proposition is similar to claim 13 of Shariff and Sheffet (2018) by showing that *any* policy satisfying the (ϵ, δ) -differential privacy in Definition 1 must suffer regret that is linear in the time horizon T . The proof of Proposition 1 is, however, different from Shariff and Sheffet (2018) because we study generalized linear models such as the logistic

regression model. We relegate the complete proof to the supplementary material.

Proposition 1. Let π be a contextual pricing policy over T periods that satisfies (ϵ, δ) -differential privacy as defined in Definition 1, with $\epsilon < \ln(2)$ and $\delta < 1/4$. Then the worst-case regret of π is lower bounded by $\Omega(T)$.

To address the challenges mentioned, Shariff and Sheffet (2018) proposed a notion of “joint DP” in the context of linear contextual bandits. We adopt this notion but refer to it as *anticipating DP*. The notion of anticipating DP highlights the key property of this definition and our focus on more general dynamic personalized pricing policies. Figure 2 gives an illustration of the anticipating differential privacy (ADP) concept. Compared with the classical differential privacy notion illustrated in Figure 1, the important difference of ADP is to restrict the output sets to prices strictly *after* a customer of interest t and to only require the distributions of *anticipating* prices (denoted by $P_{>t}$ and $Q_{>t}$) to remain stable with change of personal information $(x_t, y_t) \rightarrow (x'_t, y'_t)$ at time t . Such a restriction is motivated by the fact that the communication about (x_t, p_t, y_t) at time t is secured and the data prior to time t have no impact on the privacy of customer t because the pricing algorithm has no knowledge of x_t before time t . With the formulation of anticipating differential privacy, the challenges we mentioned earlier are resolved because the pricing decision p_t at time t is no longer in the information set of a potential attacker.

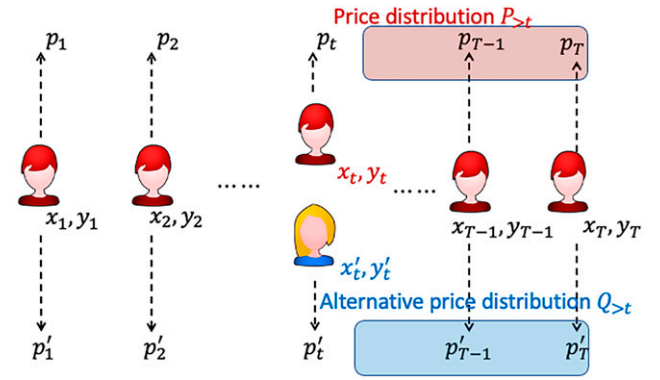
Our next definition gives a rigorous mathematical formulation of the anticipating differential privacy notion illustrated in Figure 2.

Definition 2 (Anticipating (ϵ, δ) -Differential Privacy). Let $\epsilon, \delta > 0$ be privacy parameters. A dynamic personalized pricing policy π satisfies anticipating (ϵ, δ) -differential privacy if for any pair of neighboring databases D, D' differing at time t (i.e., $o_t \neq o'_t$) and measurable set $\mathcal{P}_{>t}$, it holds that

$$\Pr[p_{t+1}, \dots, p_T \in \mathcal{P}_{>t} \mid \pi, D] \leq e^\epsilon \Pr[p_{t+1}, \dots, p_T \in \mathcal{P}_{>t} \mid \pi, D'] + \delta. \quad (3)$$

We also remark that all privacy definitions in this section are *model-free*, meaning that they do *not* depend on how realized demands y_t are modeled. Hence, the privacy guarantees of our proposed algorithm are independent from the generalized linear demand model in Equations (1) and (2). This fact is essential in practical implementations of privacy-aware algorithms because one cannot build privacy guarantees of an algorithm on a specific underlying model, which may or may not hold in reality. The modeling

Figure 2. (Color online) Illustration of the ADP Concept



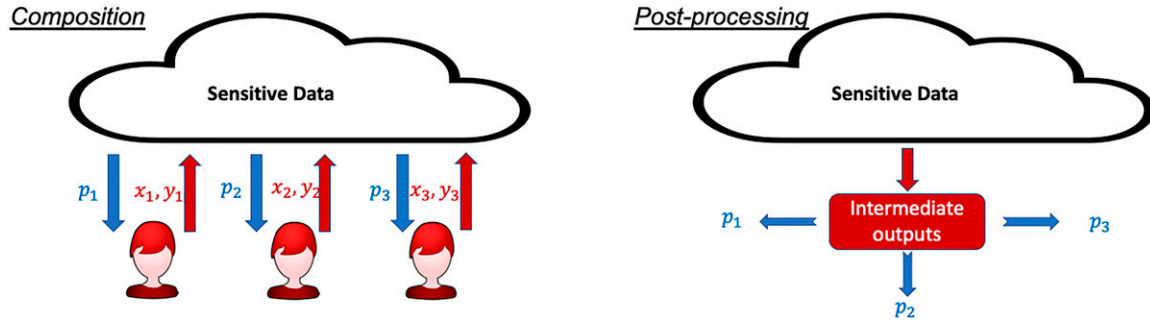
assumptions, on the other hand, are required for *performance analysis* (also known as *utility analysis*—e.g., regret upper bounds or convergence results) of our proposed privacy-aware pricing policies.

4.3. Composition in Differential Privacy

When a differentially private algorithm only outputs a single statistic (e.g., the sample mean of the database), Definition 1 is easy to check and verify. In reality, however, a useful differentially private protocol is tasked to release several statistics (sometimes with adaptively chosen queries), and the *entire* output sequence of a protocol needs to be differentially private. With multiple output statistics, Definition 1 involves high-dimensional vector spaces and is therefore difficult to check and verify. *Composition*, on the other hand, provides convenient *upper bounds* on the privacy guarantee of composite outputs using privacy guarantees of individual queries. Take the dynamic pricing setting as an example. The seller repeatedly interacts with the potential customers by offering different prices. It is therefore essential to leverage a composition guarantee in Fact 1 to make sure that all the prices offered, *when aggregated as a whole*, do not leak consumers' privacy via their personalized data.

The left panel of Figure 3 gives an illustration of the concept of composition in the context of personalized pricing. In this simple example, a centralized pricing algorithm has access to a pool of past customers' sensitive data and offers personalized prices to three customers. The rule of composition in differential privacy asserts that the privacy guarantee of the pricing algorithm *worsens* as the pricing algorithm offers prices to more customers, each time with access and calculations based on the majority of the same sensitive data. In particular, if the privacy guarantee for each individual pricing decision is ϵ , then the joint privacy guarantee when k individualized prices are offered will

Figure 3. (Color online) Illustration of the Concepts of Composition (Left) and Postprocessing (Right) in Differential Privacy



worsen to $\Omega(k\varepsilon)$ or $\Omega(\sqrt{k}\varepsilon)$, depending on the detailed composition mechanisms.

More specifically, let $A = (A_1, \dots, A_k)$ be a collection of k adaptively chosen queries, and suppose that each query A_k satisfies (ε, δ) -differential privacy as defined in Definition 1. The following result is standard in the literature and cited from theorems 3.16 and 3.20 from Dwork and Roth (2014).

Fact 1. The *composite query* $A = (A_1, \dots, A_k)$ satisfies (ε', δ') -differential privacy with either one of the following:

- (1) (Basic composition) $\varepsilon' = k\varepsilon, \delta' = k\delta$.
- (2) (Advanced composition) $\varepsilon' = \sqrt{2k \ln(1/\tilde{\delta})}\varepsilon + k\varepsilon(e^\varepsilon - 1), \delta' = k\delta + \tilde{\delta}$, for $\tilde{\delta} > 0$.

To avoid potential confusion, we remark that both basic and advanced compositions apply to *any* differentially private algorithms. Indeed, they are two different types of joint privacy guarantees proved using different techniques, reflecting different trade-offs when composing multiple differentially private queries/algorithms together. In particular, the basic composition shows a linear growth in the ε parameter (i.e., $\varepsilon' = k\varepsilon$), but it allows the δ' parameter to be 0 when the individual queries are $(\varepsilon, 0)$ -private. On the other hand, the advanced composition allows for a slower growth of the ε parameter (i.e., $\varepsilon' \asymp \sqrt{k}\varepsilon$) but must yield an (ε', δ') differential privacy guarantee with $\delta' > 0$, even if the individual queries are $(\varepsilon, 0)$ -private. In this paper, we shall use primarily the advanced composition result because we focus on (ε, δ) privacy guarantees with $\delta > 0$.

Corollary 1 (Corollary 3.21 of Dwork and Roth (2014)). *Given target privacy level $0 < \varepsilon' < 1, \delta' > 0$ of the composite query A , it is sufficient for each subquery to be (ε, δ) -differentially private with $\varepsilon = \varepsilon' / 2\sqrt{2k \ln(2k/\delta)}$ and $\delta = \delta' / 2k$.*

4.4. Postprocessing in Differential Privacy

Practical privacy-aware algorithms usually involve several separate subroutines. In most cases, not all

subroutines access the sensitive database: some subroutines may only process the results from other subroutines. The principle of *postprocessing* states that one only needs to preserve the privacy of those subroutines with access to the sensitive database in order to argue for privacy protection of the entire algorithm. For example, in dynamic pricing, algorithms are developed into different components, and only one of them directly accesses the sensitive data. It is therefore necessary to use the concept of postprocessing to argue that the entire algorithm viewed as a whole does not leak consumers' private personalized data.

The right panel of Figure 3 gives an intuitive illustration of the postprocessing concept in differential privacy. Suppose that an algorithm with full access to all sensitive data has produced some intermediate results (as shown in the red square of the illustration), and these intermediate results have already satisfied the definitions of differential privacy. Further assume that there is a downstream algorithm, which operates arbitrarily on the intermediate results to produce the personalized prices p_1, p_2, p_3, \dots , *without accessing the sensitive data any more*. Then the *postprocessing* asserts that there is no need to worry about potential privacy leakages of the downstream algorithm because the intermediate results have already been privatized. This useful concept makes it easier to design multistep, sophisticated, privacy-preserving algorithms.

More specifically, let A be a subroutine with access to the sensitive database, and let B be a subroutine that depends only on the results of A .

Fact 2 (Proposition 2.1 of Dwork and Roth (2014)). Suppose the outputs of subroutine A satisfy (ε, δ) -differential privacy. Then the outputs of subroutine B also satisfy (ε, δ) -differential privacy.

5. Algorithmic Framework

In this section we present the framework of our proposed privacy-aware dynamic personalized pricing algorithm.

A straightforward idea is to directly inject noise into customers' sensitive information (e.g., x_t) to protect privacy. However, as we will explain later in Section 9.1, such a method will fail because the features of each individual customer are relatively independent of each other. Thus, an excessively large magnitude of noise needs to be injected, which incurs a large regret. Therefore, this paper will develop a new dynamic personalized pricing algorithm based on the privacy-preserving maximum likelihood estimator. To better illustrate our algorithm, we first introduce two types of routines used in our algorithm: the *private releasers* that access the sensitive database and produce differentially private outputs and the *price optimizers* that access only the outputs from private releasers to assign near-optimal and privacy-aware prices. Then, a pseudo-code description of our main algorithm will be presented and discussed.

5.1. Private Releasers and Price Optimizers

Our proposed privacy-preserving dynamic personalized pricing algorithm consists of several subroutines. We divide the subroutines into two classes: the private releasers and the price optimizers.

The *private releasers* access the sensitive database $\{x_t, p_t, y_t\}_{t=1}^T$ and output differentially private intermediate results. For example, in Figure 4, the PRIVATECOV routine returns differentially private sample covariance matrices, and the PRIVATEMLE routine returns differentially private maximum likelihood estimates. For private releaser routines, the differential privacy notions are classical (in Definition 1). Note that, in addition to differential privacy guarantees, the subroutines also need to satisfy the anticipating constraints for pricing algorithms (i.e., accessing only $\{x_\tau, y_\tau, p_\tau\}_{\tau < t}$ to produce any outputs being used at time t).

The *price optimizer*, on the other hand, performs optimization and outputs the prices p_t for each time period t . To ensure privacy, our designed price optimizer will *not* directly access historical sensitive data $\{x_\tau, y_\tau, p_\tau\}_{\tau < t}$. Instead, it optimizes the offering price p_t based only on x_t (the personal information of the incoming customer) and intermediate quantities computed by private releasers up to time t .

Because our designed price optimizer has access to x_t at time t , one cannot directly apply the postprocessing rule in Fact 2 to argue privacy guarantees. Nevertheless, the following proposition shows that if all private releasers are differentially private, then so is the price optimizer in the sense of anticipating differential privacy in Definition 2. The proof of Proposition 2 is placed in the supplementary material.

Proposition 2. *Let (a_1, \dots, a_T) be the outputs of private releasers at each time period t , and suppose the entire output sequence (a_1, \dots, a_T) satisfies (ϵ, δ) -differential privacy.*

Suppose the price p_t at time t is a deterministic function of x_t and a_1, \dots, a_{t-1} . Then the pricing policy satisfies anticipating (ϵ, δ) -differential privacy.

Remark 1. The conclusion in Proposition 2 holds for p_t as randomized functions of x_t, a_1, \dots, a_{t-1} as well. Nevertheless, because in our proposed algorithm the price optimizer is deterministic, we shall restrict ourselves to deterministic functions.

5.2. Our Policy

In Figure 4 we depict a high-level framework of our privacy-aware dynamic personalized pricing policy. It shows a three-layer structure of the proposed policy. The first layer is the *sensitive database*, consisting of data $\{o_t = (p_t, x_t, y_t)\}_{t=1}^T$; its privacy needs to be protected. The second layer comprises *private releasers*, consisting of two subroutines, PRIVATECOV (see Algorithm 2 in Section 6.1) and PRIVATEMLE (see Algorithm 3 in Section 6.2). The PRIVATECOV subroutine supplies differentially private sample covariance matrices $\Lambda_n^p \in \mathbb{R}^{d \times d}$ at every time period. The PRIVATEMLE subroutine outputs differentially private maximum likelihood estimates $\hat{\theta}_n^p$ —but only when such estimates are requested by the price optimizer. The PRIVATECOV subroutine is designed to be (ϵ_1, δ_1) -differentially private, and the PRIVATEMLE routine is (ϵ_2, δ_2) -differentially private, so that all outputs from private releasers are $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differentially private, thanks to the basic composition rule in Fact 1.

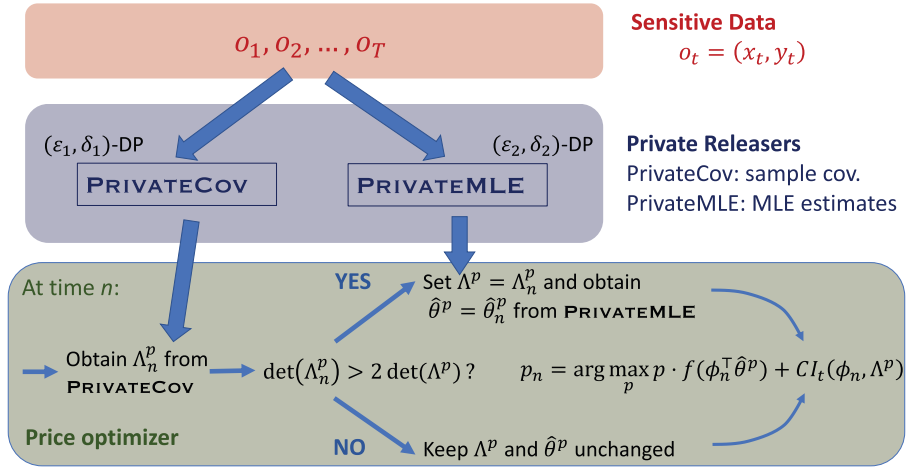
The third layer of our proposed policy is the price optimizer. As discussed in the previous section, to ensure privacy, the price optimizer shall not access the sensitive database D directly. Instead, it should base its decision of p_t on outputs from private releasers and x_t only. The last block in Figure 4 illustrates the basic flow of our price optimizer. The price optimizer maintains Λ^p and $\hat{\theta}^p$ throughout the pricing process, both of which are obtained directly from private releasers without accessing the sensitive database. At the beginning of time period n , the price optimizer first obtains sample covariance Λ_n^p from the PRIVATECOV routine. The optimizer then decides whether to request fresh MLE from the PRIVATEMLE routine by comparing $\det(\Lambda_n^p)$ with $\det(\Lambda^p)$, in addition to some other criteria specified in Algorithm 1. Afterward, p_t is selected as the maximizer of an upper confidence bound of the expected revenue on x_t . It is only during this step that the personal information x_t is involved.

Algorithm 1 (The framework of privacy-aware dynamic personalized pricing)

1: **Input:** privacy parameters $\epsilon_1, \delta_1, \epsilon_2, \delta_2 > 0$, number of pure-exploration periods T_0 , maximum number of PRIVATEMLE calls D_∞ , regularization parameter $\rho \geq 1$, confidence parameter $\gamma > 0$.

2: **Output:** the offering prices p_1, p_2, \dots, p_T ;

Figure 4. (Color online) Our Algorithm Framework



Note. Details and explanations in Section 5.2 in the main text.

- 3: $\delta_2' = \frac{\delta_2}{2D_\infty}$, $\varepsilon_2' \leftarrow \frac{\varepsilon_2}{2\sqrt{2D_\infty \ln(1/\delta_2')}}}$, $\Lambda^p = \rho I_d$, $\hat{\theta}^p = 0$, $D_{MLE} = 0$;
- 4: For the first T_0 time periods, offer prices p_t uniformly at random from $[0, 1]$;
- 5: **for** $n = T_0 + 1, \dots, T$ **do**
- 6: Obtain $\Sigma_n^p \leftarrow \text{PrivateCov}(n, \varepsilon_1, \delta_1)$ and let $\Lambda_n^p = \Sigma_n^p + \rho I_d$;
- 7: **if** $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ and $D_{MLE} < D_\infty$ **then**
- 8: $\hat{\theta}^p \leftarrow \text{PrivateMLE}(n, \rho, \varepsilon_2', \delta_2')$, $\Lambda^p \leftarrow \Lambda_n^p$, $D_{MLE} \leftarrow D_{MLE} + 1$;
- 9: **end if**
- 10: Offer price $p_n = \arg \max_{p \in [0, 1]} \min \{1, pf(\phi_n^T \hat{\theta}^p) + \gamma \sqrt{\phi_n^T (\Lambda^p)^{-1} \phi_n}\}$, where $\phi_n = \phi(x_n, p_n)$;
- 11: **end for**

Algorithm 1 also gives a pseudo-code description of our proposed pricing policy, which is more accurate and detailed than Figure 4. Note that Algorithm 1 involves several algorithmic parameters, such as T_0, D_∞, γ , and ρ , which do *not* affect the privacy guarantees of the algorithm but do have an impact on its performance. How to set these algorithmic parameters will be given later in Section 7, when we analyze the regret performance of Algorithm 1. Before that, we will first make a few important remarks about Algorithm 1.

Remark 2 (Time Complexity).

The time complexity for the PRIVATEMLE subroutine is the same as traditional MLE calculations, if not easier (because the overall formulation is convex), because only the objective is perturbed with a linear term. The time complexity for the PRIVATECOV subroutine is slightly more expensive: at each time n , the tree-based protocol needs to update $O(\log n)$ nodes on the binary tree instead of just adding $\phi_t \phi_t^T$ to a counting matrix. Overall, the algorithm's time complexity is $O(d^3 T \ln T)$

(note d^3 comes from the computation of the determinant), in addition to $O(d \ln T)$ number of MLE calculations. The next section gives more details on the two private releasers.

Remark 3 (Difference from Generalized Linear Contextual Bandit).

This remark explains how the algorithm differs from a classic generalized linear bandit algorithm without privacy consideration. The major difference is that when there is no privacy consideration, there is no need (and no use) to randomize and therefore vanilla MLE can be used to obtain an estimated model $\hat{\theta}_t$ at every time period t , with standard statistical analysis of the errors for such estimates (see Li et al. (2017)). With privacy constraints, such maximum likelihood estimates need to be carefully privatized by calibrating artificial noise into the objective of the MLE (the PRIVATEMLE subroutine later in Algorithm 3), which also calls for more detailed perturbation-based statistical analysis. Another difference is that without privacy constraints, the seller could update its model estimate $\hat{\theta}_t$ at every time period to obtain the most accurate and updated information. With privacy constraints, however, the seller cannot afford to adaptively compute a model estimate after each time period as a result of composition constraints and must perform such model estimates sparingly, relying further on a signal scheme also privatized by incorporating artificial noise matrices (see the PRIVATECOV subroutine later in Algorithm 2).

Remark 4 (Exploration Phase).

In addition, we also clarify that the forced exploration step in our algorithm is optional: the proposed algorithm remains valid (i.e., satisfying suitable differential privacy constraints and achieving small overall regret) without the forced exploration step (see

Theorem 1 in Section 7.1, where $T_0 = 0$). The forced exploration helps to ensure improved regret guarantee when there are additional distributional assumptions on contextual vectors (see Section 7.2). This forced exploration aims to make sure the sample covariance of the context vectors is well conditioned, which leads to improved regret guarantees of privatized MLE.

6. Design and Analysis of Private Releasers

In this section, we give detailed designs of the two private releasers: the `PRIVATECOV` subroutine and the `PRIVATEMLE` subroutine. We prove that both of them satisfy (ϵ, δ) -differential privacy as defined in Definition 1. We also prove several utility guarantees that will be helpful later in the regret analysis of the pricing policy. Figure 5 shows the flow of our proof framework. Because of space constraints and exposition concerns, all proofs to technical lemmas or propositions in this section have been placed in the supplementary material.

6.1. The `PRIVATECOV` Subroutine

Algorithm 2 gives a pseudo-code description of the `PRIVATECOV` subroutine. Note that in Algorithm 2 the Σ_n^p covariance matrices are released sequentially once each time period, and `PrivateCov`(n, ϵ, δ) would simply be the Σ_n^p matrix released at the end of iteration $n - 1$.

Algorithm 2 is based on the `AnalyzeGauss` framework in Dwork et al. (2014) coupled with the *tree-based aggregation* technique for releasing continual observations (Dwork et al. 2010, Chan et al. 2011). The `AnalyzeGauss` method by Dwork et al. (2014) develops a Gaussian mechanism on releasing a single covariance matrix privately from the data. On the other hand, tree-based aggregation provides a general protocol on how to continually release sequentially updated statistics (e.g., partial sums of sample covariance matrices) under privacy constraints. For our `PRIVATECOV`, by calibrating

symmetric random Gaussian matrices $\{W^n\}$ into the sample covariances under the tree-based aggregation, one achieves differential privacy. The following proposition claims that the outputs $(\Sigma_1^p, \dots, \Sigma_{T-1}^p)$ of Algorithm 2 satisfy (ϵ, δ) -differential privacy.

Algorithm 2 (The `PRIVATECOV` subroutine)

- 1: **function** `PRIVATECOV`(T, ϵ, δ) \triangleright returns $\Sigma_1^p, \dots, \Sigma_{T-1}^p$
- 2: $\delta' \leftarrow \frac{\delta}{2\lceil \log_2 T \rceil}, \quad \epsilon' \leftarrow \frac{\epsilon}{2\lceil \log_2 T \rceil \ln(1/\delta')}, \quad \sigma_{\epsilon', \delta'}^2 = \frac{2\ln(1.25/\delta')}{(\epsilon')^2}, \quad m = \lceil \log_2 T \rceil;$
- 3: **Initialize** $\Sigma(\ell) = \hat{\Sigma}(\ell) = 0$ for all $\ell = 0, \dots, m - 1;$
- 4: **for** $n = 1, 2, \dots, T - 1$ **do**
- 5: Express n in its binary form: $n = \sum_{\ell=0}^{m-1} b_n(\ell)2^\ell, b_n(\ell) \in \{0, 1\};$
- 6: Let $\ell_n \leftarrow \min \{\ell : b_n(\ell) = 1\}$ be the least significant bit of $n;$
- 7: Update $\Sigma(\ell_n) \leftarrow \phi_n \phi_n^\top + \sum_{\ell < \ell_n} \Sigma(\ell)$ and $\hat{\Sigma}(\ell) \leftarrow \hat{\Sigma}(\ell) \leftarrow 0$ for all $\ell < \ell_n;$
- 8: Calibrate noise: $\hat{\Sigma}(\ell_n) \leftarrow \Sigma(\ell_n) + W^n,$ where $W_{ij}^n = W_{ji}^{n.i.d.} \sim \mathcal{N}(0, \sigma_{\epsilon', \delta'}^2);$
- 9: Release $\Sigma_n^p = \sum_{\ell=0}^{m-1} b_n(\ell) \hat{\Sigma}(\ell);$
- 10: **end for**
- 11: **end function**

Proposition 3. The outputs of Algorithm 2, $(\Sigma_1^p, \dots, \Sigma_{T-1}^p)$, satisfy (ϵ, δ) -differential privacy.

The following lemma further gives high-probability bounds on the deviation from Σ_n^p to the actual sample covariance $\Sigma_n = \sum_{t=1}^n \phi_t \phi_t^\top$. This utility guarantee is useful later in the regret analysis to justify the $\det(\Lambda_n^p) > 2\det(\Lambda^p)$ condition in Algorithm 1.

Lemma 1. With probability $1 - O(T^{-1})$, it holds for all $n \in \{1, 2, \dots, T - 1\}$ that

$$\|\Sigma_n^p - \Sigma_n\|_{\text{op}} \leq O(\epsilon^{-1} \sqrt{d} \ln^{4.5}(T/\delta)),$$

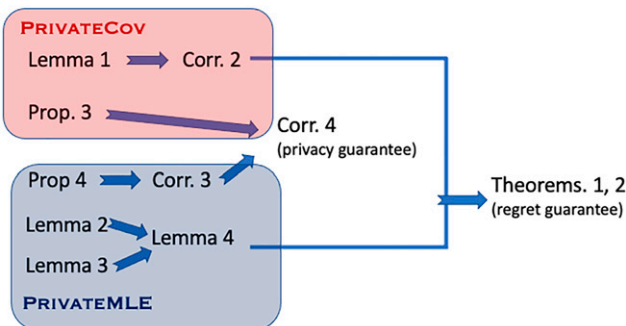
where $\Sigma_n = \sum_{t \leq n} \phi_t \phi_t^\top$.

The following corollary is an immediate consequence of Lemma 1.

Corollary 2. Let $\Lambda_n = \Sigma_n + \rho I_d$ and $\Lambda_n^p = \Sigma_n^p + \rho I_d$ for some $\rho \geq \epsilon^{-1} d \sqrt{d} \ln^5(T/\delta)$. Then there exists a universal constant $C_T < \infty$ such that for any $T \geq C_T$, with probability $1 - O(T^{-1})$ for all $n \in \{1, 2, \dots, T - 1\}$, it holds that $0.9\det(\Lambda_n) \leq \det(\Lambda_n^p) \leq 1.11\det(\Lambda_n)$.

Corollary 2 shows that when the `PrivateMLE` is invoked in Algorithm 1, the determinant of the real sample covariance matrix roughly doubles. This is important to our later regret analysis because if `PRIVATEMLE` is invoked too frequently, the algorithm pays the price of composition in privacy. Although, on the other hand, if `PRIVATEMLE` is invoked too rarely, the old (and inaccurate) parameters will be used for a long time, which incurs a larger regret. Therefore, our

Figure 5. (Color online) Flow of Our Proof Framework



Note. Corr., corollary; Prop., proposition.

analysis shows that the right frequency should be invoking PRIVATEMLE once the determinant of the privacy-preserving covariance roughly doubles.

6.2. The PRIVATEMLE Subroutine

Algorithm 3 gives a pseudo-code description of the PRIVATEMLE subroutine. The algorithm is based on the objective perturbation framework developed in Chaudhuri et al. (2011) and Kifer et al. (2012). More specifically, Algorithm 3 calibrates a noisy term ($w^\top \theta$) into the constrained MLE formulation in order to achieve differential privacy of the output optimal solutions $\hat{\theta}_n^p$.

The following proposition establishes the claim that Algorithm 3 is (ϵ, δ) -differentially private.

Proposition 4. *The output of Algorithm 3, $\hat{\theta}_n^p$, satisfies (ϵ, δ) -differential privacy.*

The next corollary, which establishes the privacy guarantee of the PRIVATEMLE, immediately follows Proposition 4 and Corollary 1. It shows how to set the algorithmic parameters in Algorithm 3 to ensure that the resulting price decisions are differentially private at the designated levels ϵ and δ .

Corollary 3. *Suppose PRIVATEMLE is invoked for at most D_∞ times in Algorithm 1. Then the composite sequence of D_∞ outputs of PRIVATEMLE satisfies (ϵ, δ) -differential privacy if each call of PRIVATEMLE is supplied with privacy parameters $\delta' = \frac{\delta}{2D_\infty}$ and $\epsilon' = \frac{\epsilon}{\sqrt{2D_\infty \ln(1/\delta')}}$.*

Now, we are ready to provide the privacy guarantee of the entire policy in Algorithm 1.

Corollary 4. *The price decisions $\{p_1, \dots, p_T\}$ of Algorithm 1 satisfy $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differential privacy.*

Algorithm 3 (The PRIVATEMLE subroutine)

- 1: **function** PRIVATEMLE($n, \rho, \epsilon, \delta$) \triangleright returns $\hat{\theta}_n^p$
- 2: $B_1 \leftarrow (B_Y + 1)G, B_2 \leftarrow KG, \rho \leftarrow \max\{\rho, 2B_2/\epsilon\}, v_{\epsilon, \delta}^2 \leftarrow B_1^2(8\ln(2/\delta) + 4\epsilon)/\epsilon^2;$
- 3: Sample $w \sim \mathcal{N}(0, v_{\epsilon, \delta}^2 I_d);$
- 4: Return $\hat{\theta}_n^p = \arg \min_{\|\theta\|_2 \leq 2} \left\{ \left(\sum_{t < n} t < n - \ln p(y_t | \phi_t, \theta) \right) + \frac{\rho}{2} \|\theta\|_2^2 + w^\top \theta \right\};$
- 5: **end function**

Corollary 4 immediately follows Proposition 3, Corollary 3, Proposition 2, and Fact 1. More specifically, in Algorithm 1, PRIVATECOV is invoked with parameters (ϵ_1, δ_1) , which is (ϵ_1, δ_1) -differential privacy. Moreover, PRIVATEMLE is invoked with parameters (ϵ_2', δ_2') for at most D_∞ times, whose outputs are (ϵ_2, δ_2) -differential privacy. Therefore, the entire policy satisfies $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -differential privacy, thanks to Proposition 2 and the basic composition rule in Fact 1.

In the rest of this section we establish the prediction error guarantee (a.k.a. the utility guarantee in

differential privacy literature) of the estimator $\hat{\theta}_n^p$ from the PRIVATEMLE subroutine. More precisely, Lemma 2 upper bounds the prediction errors of the sequence of obtained model estimate $\hat{\theta}_n^p$ with the presence of artificially calibrated noises in the PRIVATEMLE subroutine. With smaller values of ϵ, δ indicating stronger privacy protection, the parameter $v_{\epsilon, \delta}$ becomes larger, which leads to a larger variance of the Gaussian noise w . Thus, the PRIVATEMLE subroutine needs to calibrate higher magnitudes of noise into the objective function, leading to either larger prediction errors (see (4)) or a lengthened forced exploration phase (see the condition of (5)). Note that the lower bound on $\lambda_{\min}(\Sigma_n)$ for (5) is achieved by the exploration phase in Algorithm 1. This key result quantifies the trade-off between the strength of the privacy protection and the prediction errors of the model parameter estimates. Another point of practical guidance provided by Lemma 2 is that the regularization amount ρ also needs to grow as ϵ, δ becomes smaller.

We emphasize that this key utility guarantee in Lemma 2 is *not* directly covered by existing utility analysis in Chaudhuri et al. (2011) or Kifer et al. (2012) for two reasons. First, in Chaudhuri et al. (2011) and Kifer et al. (2012), the utility is measured in terms of the difference between *objective values* before and after objective perturbation, which is *not* sufficient for the purpose of analyzing contextual bandit algorithms that require first-order KKT conditions. Additionally, in both Chaudhuri et al. (2011) and Kifer et al. (2012), the data (ϕ_t, y_t) are assumed to be sampled *independently and identically* from an underlying distribution, whereas in our problem, the data clearly are neither independent nor identically distributed.

We also remark that our utility analysis of the (differentially private) constrained MLE (see the proof of Lemma 2) differs significantly from existing analysis of generalized linear contextual bandit problems as well (Filippi et al. 2010, Li et al. 2017, Wang et al. 2019). In Li et al. (2017), it is assumed that ϕ_t are i.i.d. and their distributions satisfy a certain nondegenerate assumption, which we do not necessarily impose in this paper. In both Filippi et al. (2010) and Wang et al. (2019), the formulations of the optimization problems are nonconvex in θ , which facilitates the analysis of the properties of the optimal solution. However, the nonconvex formulation poses significant challenges for privacy-aware algorithms because differentially private methods for nonconvex optimization are scarce. It is therefore a highly nontrivial task to analyze a fully convex optimization formulation without stochasticity assumptions on ϕ_t .

Lemma 2. Fix $n \in \{1, 2, \dots, T-1\}$, and let $\Lambda_n = \Sigma_n + \rho I = \sum_{t < n} \phi_t \phi_t^\top + \rho I$. Suppose $\rho \geq \max\{5v_{\epsilon, \delta} \sqrt{5d \ln T}, 2 + 48s^2 G^2 K d \ln T\}$. Then with

probability $1 - O(T^{-2})$, the following hold: $\|\hat{\theta}_n^p\|_2 < 2$, and

$$(\hat{\theta}_n^p - \theta^*)^T \Lambda_n (\hat{\theta}_n^p - \theta^*) \leq (sK\sqrt{3d\ln T} + (2G+3)\sqrt{\rho} + Gv_{\varepsilon,\delta}\sqrt{5d\ln T})^2. \quad (4)$$

Furthermore, if $\lambda_{\min}(\Sigma_n) \geq \lambda_0 = \left[\frac{(2G+3)\rho}{\sqrt{5d\ln T}} + v_{\varepsilon,\delta}G\right]^2$, then Inequality (4) can be strengthened to

$$(\hat{\theta}_n^p - \theta^*)^T \Lambda_n (\hat{\theta}_n^p - \theta^*) \leq [4sK\sqrt{d\ln T}]^2. \quad (5)$$

Lemma 2 is proved by analyzing the first-order KKT condition at $\hat{\theta}_n^p$ and is deferred to the supplementary material. Lemma 2 upper bounds the transformed estimation error of the differentially private MLE $\hat{\theta}_n^p$ in two upper bounds. The first upper bound in (4) applies to the general setting and has a $Gv_{\varepsilon,\delta}\sqrt{5d\ln T}$ additive term involving the differential privacy parameters ε and δ in the upper bound. The second upper bound in (5), on the other hand, shows that if the sample covariance matrix Σ_n is spectrally lower bounded, then the upper bound on $\|\hat{\theta}_n^p - \theta^*\|_{\Lambda_n}$ can be much improved with only the standard $O(\sqrt{d\ln T})$ term.

7. Regret Analysis

Section 6 has established the privacy guarantees of our dynamic personalized pricing policy (see Corollary 4). In this section, we will further analyze the performance/utility of our proposed policy by proving upper bounds on its expected cumulative regret.

Recall that in the dynamic personalized pricing problem, there are t time periods, and at each time period, a customer arrives with personal information x_t . When offered price p_t , the expected demand is modeled by the generalized linear model $p(y_t | p_t, x_t, \theta^*) = \exp\{\zeta(y\phi(p_t, x_t)^T \theta^* - m(\phi(p_t, x_t)^T \theta^*) + h(y_t, \zeta))\}$ with expectation $\mathbb{E}[y_t | p_t, x_t, \theta^*] = f(\phi(p_t, x_t)^T \theta^*)$. With θ^* known in hindsight, the optimal price p_t^* at time t is the one maximizing the retailer's expected revenue, or more specifically,

$$p_t^* := \arg \max_{p \in [0,1]} pf(\phi(p, x_t)^T \theta^*).$$

The regret of a dynamic pricing policy π is then defined as the cumulative difference between the expected revenue of the policy's offered prices and that of a clairvoyant, or more specifically,

$$\text{Regret}(\pi; T) := \sum_{t=1}^T p_t^* f(\phi(p_t^*, x_t)^T \theta^*) - p_t f(\phi(p_t, x_t)^T \theta^*).$$

Clearly, by definition, the regret of any admissible policy is always nonnegative because no p_t has a higher expected revenue compared with p_t^* . The smaller

the regret, the better the policy's performance. We are also primarily focused on the asymptotic growth of the regret as a function of the time horizon T , as well as several other important parameters, such as the feature dimension d and the privacy parameters $\varepsilon_0 := \varepsilon_1 + \varepsilon_2$, $\delta_0 := \delta_1 + \delta_2$.

7.1. The General Case

We first analyze the regret of Algorithm 1 in the most general case, in which customers' personal information $\{x_t\}$ is obviously (i.e., prefixed) but can be adversarially chosen without preassumed patterns. Our next theorem upper bounds the regret of Algorithm 1 with proper choices of the values of algorithmic parameters. Recall that $\varepsilon_0 := \varepsilon_1 + \varepsilon_2$, $\delta_0 := \delta_1 + \delta_2$. We also note that for the general case, the random exploration phase (step 4 in Algorithm 1) will be unnecessary, and thus we could set $T_0 = 0$.

Theorem 1. Suppose Algorithm 1 is run with parameters $\varepsilon_1, \varepsilon_2 \geq 0.1\varepsilon_0$, $\delta_1, \delta_2 \geq 0.1\delta_0$, $T_0 = 0$, $D_\infty = \lceil d \log_{1.5} T \rceil$, $\rho = \max\{\varepsilon_1^{-1} d^{1.5} \ln^5 T, 5v_{\varepsilon_2', \delta_2'} \sqrt{5d\ln T}, 2 + \frac{48s^2 G^2}{Kd\ln T}\}$, and $\gamma = K[(\sqrt{3sK} + \sqrt{5Gv_{\varepsilon_2', \delta_2'}}) \sqrt{d\ln T} + (2G+3)\sqrt{\rho}]$, where $\varepsilon_2', \delta_2'$ are defined in step 3 of Algorithm 1 and $v_{\varepsilon_2', \delta_2'}$ is defined in Algorithm 3. Then it holds that

$$\text{Regret}(\pi; T) \leq 2\gamma\sqrt{4.6d\ln T} \leq \tilde{O}\left(\varepsilon_0^{-1}\sqrt{d^3 T \ln^5(1/\delta_0)}\right),$$

where in the $\tilde{O}(\cdot)$ notation, we omit logarithmic terms in T and polynomial dependency on other model parameters s , K , G , and B_Y .

Theorem 1 is proved in the supplementary material. We note that when T is large, our regret bound matches the classical optimal regret bound of $O(\sqrt{T})$. The dependency on the dimensionality of personal information d (i.e., $\sqrt{d^3}$) can be further improved by assuming a stronger assumption on the stochasticity of personal information x_t (see Section 7.2). The stochastic personal information or demand covariate has been a common assumption in the pricing literature (see, e.g., Qiang and Bayati (2016), Javanmard and Nazerzadeh (2019), Ban and Keskin (2021), and Chen et al. (2021)).

7.2. Improved Regret with Stochastic Contexts

In this section, we show that for a large class of problems in which the customers' personal information is stochastically distributed, the regret upper bound in Theorem 1 could be significantly sharpened.

The following assumption mathematically characterizes the stochasticity condition of customers' personal information used in this section.

Assumption 1. Let $U[0,1]$ be the uniform distribution on $[0,1]$. There exists an underlying distribution μ_x and a constant $\kappa_x > 0$ such that $x_1, \dots, x_T \stackrel{i.i.d.}{\sim} \mu_x$, and

furthermore,

$$\begin{aligned} \|\phi(x, p)\|_2 &\leq 1 \text{ a.s. } \sim \mu_x \times U[0, 1]; \\ \mathbb{E}_{(x,p) \sim \mu_x \times U[0,1]}[\phi(p, x)\phi(p, x)^\top] &\geq \kappa_x I_d. \end{aligned}$$

Assumption 1 assumes that consumers' personal feature vectors are relatively widely spread so that they are not concentrated in a narrow region or direction. Such an assumption helps improve the regret analysis because the algorithm can expect to see feature vectors along with any directions with reasonable chances, and therefore the overall estimates of the unknown regression model can be more accurate.

With Assumption 1, the following theorem shows that when algorithmic parameters are properly chosen in Algorithm 1, the regret upper bound can be improved compared with Theorem 1 for the general setting.

Theorem 2. Under Assumption 1, suppose Algorithm 1 is run with parameters $\varepsilon_1, \varepsilon_2 \geq 0.1\varepsilon_0$, $\delta_1, \delta_2 \geq 0.1\delta_0$, $D_\infty = \lceil d \log_{1.5} T \rceil$, $\rho = \max\{\varepsilon_1^{-1} d^{1.5} \ln^5 T, 5v_{\varepsilon_2', \delta_2'} \sqrt{5d \ln T}, 2 + 48\varepsilon_2^2 G K d \ln T\}$, $T_0 = 32 \lceil \frac{(2G+3)\rho}{\sqrt{5d \ln T}} + v_{\varepsilon, \delta} G \rceil^2 \ln^2(dT)$, $\gamma = 4sK^2 \sqrt{d \ln T}$, where $\varepsilon_2', \delta_2'$ are defined in step 3 of Algorithm 1 and $v_{\varepsilon_2', \delta_2'}$ is defined in Algorithm 3. Then it holds for sufficiently large $T \geq e^{\kappa_x^2}$ that

$$\begin{aligned} \text{Regret}(\pi, T) &\leq T_0 + 2\gamma \sqrt{4.6dT \ln T} \\ &\leq \tilde{O}\left(d\sqrt{T} + \varepsilon_0^{-2} d^2 \ln^{10}(1/\delta_0)\right), \end{aligned}$$

where in the $\tilde{O}(\cdot)$ notation we omit logarithmic terms in T and polynomial dependency on other model parameters s, K, G , and B_Y .

The proof of Theorem 2 is largely the same as the proof of Theorem 1 except for the application of the second upper bound in Lemma 2. We relegate the complete proof of Theorem 2 to the supplementary material. Comparing Theorem 2 with Theorem 1, we note that the significant improvement lies in the additive nature between the ε_0, δ_0 and d, T terms in Theorem 2. More specifically, because the privacy-incurred terms are now additive and do not scale polynomially with T , in most practical scenarios when the time horizon T is very large, the dominating term of Theorem 2 becomes only $\tilde{O}(d\sqrt{T})$, which is optimal (up to logarithmic factors) in both the time horizon T and the feature dimension d (see, e.g., the $\Omega(d\sqrt{T})$ lower bound in Dani et al. (2008)).

7.3. Impact of Privacy Constraints on Seller Surplus

In our theoretical framework, the seller surplus is measured and reflected by the notion of *regret*, which measures how much revenue/profits are lost by the

seller's pricing decisions compared with the optimal personalized prices in hindsight. The smaller the regret, the larger the seller surplus.

Our main results in Theorems 1 and 2 give quantitative upper bounds on the regret of our proposed algorithm. More specifically, the regret of our algorithm is (omitting logarithmic factors and secondary model parameters) $\tilde{O}(\varepsilon^{-1} \sqrt{d^3 T})$ in the general setting and $\tilde{O}(\sqrt{d^2 T} + \varepsilon^{-2} d^2)$ with additional assumptions on the distribution of consumers' context vectors. Here, T is the time horizon (i.e., the number of customers handled), d is the number of covariates in consumers' personal data, and $\varepsilon > 0$ dictates the level of privacy leakage, with a smaller ε indicating stronger/stricter protection of users' privacy. From these results, we make the following observations.

- *Trade-offs between seller profits and privacy protection:* With stronger privacy protection (i.e., $\varepsilon \rightarrow 0^+$), it is clear that the regret of our proposed algorithm increases, indicating that the seller profits are going to suffer with additional privacy constraints. The decrease of seller surplus is, however, alleviated when the consumers' context vectors are relatively well distributed, as the $\varepsilon^{-2} d^2$ term is *not* the dominating term in the regret bound when there are a sufficient number of customers/users. Such a decrease of seller profits is intuitive and expected, because additional privacy constraints limit sellers' ability to offer very personally tailored prices to boost their revenues.

- *Value and privacy costs of information:* The d parameter in the regret bound characterizes how many covariates or factors the pricing algorithm exploits in customers' personalized data and shows the value and privacy costs of information: with more factors/covariates (i.e., larger values of d), the retailer is able to consider more refined details and information of each incoming customer, but such information adds to the burden of privacy protection, leading to increased regret. To see this more clearly, with some stochasticity assumption of covariates, the regret bound $\tilde{O}(d\sqrt{T} + \varepsilon^{-2} d^2)$ in Theorem 2 shows the following fact. For the regret term $\varepsilon^{-2} d^2$ related to the privacy to be a constant, a larger dimension d (i.e., more customer information) implies that $\varepsilon = C_0 d$ also grows proportionally, which leads to a weaker privacy protection. Additionally, for the first term $\tilde{O}(d\sqrt{T})$, there is also a known lower bound showing that any policy must suffer a regret of $\Omega(d\sqrt{T})$ in the worst case (Dani et al. 2008). Therefore, there is indeed a cost of information for the purpose of privacy protection. Our regret upper bounds therefore provide in principle a bottom line for practitioners to gauge the costs of incorporating more factors of user information into dynamic personalized price decisions.

8. Numerical Results

In this section we corroborate the theoretical guarantees established in this paper for our proposed

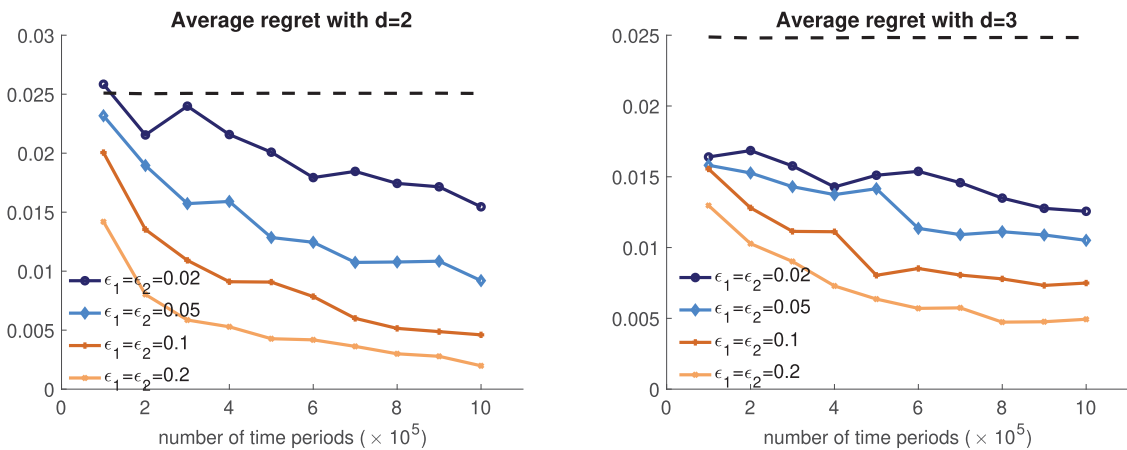
differentially private personalized pricing method with simulation results on a synthetic data set. We adopt the logistic regression model $\Pr[y_t = 1 | \phi_t, \theta^*] = \frac{e^{\zeta \phi_t^T \theta^*}}{1 + e^{\zeta \phi_t^T \theta^*}}$, with $\zeta = 4$, $\phi_t(x_t, p_t) = \frac{1}{\sqrt{d}}[x_t; -p_t] \in \mathbb{R}^d$, and $\theta^* = [-\sqrt{0.1}; -\sqrt{0.1}; \dots; -\sqrt{0.1}; \sqrt{1 - 0.1(d-1)}] \in \mathbb{R}^d$. The personal feature vectors $\{x_t\}$ are synthesized uniformly at random from the unit cube $[-1, 1]^{d-1}$. It is easy to verify that $\|\phi_t\|_2 \leq 1$ and $\|\theta^*\|_2 \leq 1$ always hold for all d . Algorithm parameters (as inputs in Algorithm 1) are chosen as $T_0 = 10$, $\rho = 10$, $D_\infty = \lceil d \log_2 T \rceil$, and $\gamma = 1$. Other privacy-related parameters will be varied to demonstrate a spectrum of our proposed algorithm on a continuous landscape of differential privacy guarantees. Note that this experiment's main purpose is to investigate the impact of privacy-related parameters (i.e., ϵ and δ) rather than compete with state-of-the-art nonprivate pricing algorithms.

In Figure 6 we plot the average regret of our proposed algorithm under various ϵ_1, ϵ_2 privacy settings and time horizons T ranging from 10^5 to 10^6 . All settings are run for 20 independent trials, and the average regret is reported. For reference purposes, we also indicate in both plots of Figure 6 (see the flat dashed line) the average regret of a policy that simply produces uniformly at random prices p_t at each t , completely ignoring the personalized features/factors of each incoming customer. As we can see, under most privacy settings, including highly secured settings with small ϵ (e.g., $\epsilon_1 = \epsilon_2 = 0.02$), the average regret of our proposed algorithm is much smaller compared with completely random prices, demonstrating its utility under privacy constraints. Furthermore, with relaxed privacy requirements (i.e., larger values of ϵ_1, ϵ_2) and/or longer pricing horizons T , the average regret of our algorithm significantly decreases, which verifies the theoretical regret upper bounds we established in Theorems 1 and 2.

In Figures 7 and 8, we provide some additional auxiliary simulation results. Figure 7 gives a direct landscape of the average regret of our algorithm under ϵ values ranging from 0.1 to 1. Figure 8 further explores the robustness of our algorithm under several very small δ values (as small as $\delta = 1/T^{10}$). Note that in Figure 8 there are multiple trend lines corresponding to the performances of the proposed algorithm under different settings of T, ϵ , and δ values. Apart from the dependency on $\ln(1/\delta)$, Figure 8 also shows that the average regret of our algorithm decreases with increasing time horizon T and relaxed privacy guarantees (i.e., larger values of ϵ), both of which are consistent with the findings in Figures 6 and 7. The results in both figures are as expected (significant decreases in average regret with large ϵ values and moderate increases in average regret with geometrically decreasing δ values) from our theoretical results.


To better illustrate our algorithm, we further report two additional sets of simulation results. In Table 1 we report the average regret of our proposed algorithm together with an algorithm that is not subject to any kind of privacy constraints, which is implemented by removing all noise calibration steps in the two private releasers PRIVATECOV and PRIVATEMLE. We remark that even larger ϵ values (e.g., $\epsilon = 1.0$ or $\epsilon = 5.0$) indicate quite nontrivial universal privacy protection of consumers' sensitive data, which explains the relatively larger regret incurred by differentially private pricing algorithms compared with their nonprivate counterparts. In Table 2 we report the average regret of our proposed algorithm when the values of ϵ_1 and ϵ_2 are very different to see which privacy parameter has a bigger impact on the performance of the designed algorithm. Table 2 shows that ϵ_2 clearly has a much larger effect on the regret performance of our algorithm, with the average regret significantly

Figure 6. (Color online) Average Regret of Our Proposed Algorithm Under Different Time Horizons T



Notes. The black dashed line indicates the average regret of a policy offering completely at random prices. Both δ_1, δ_2 parameters are set at $\delta_1 = \delta_2 = 1/T^2$.

Figure 7. (Color online) Average Regret of Our Proposed Algorithm Under Different Privacy Parameters $\varepsilon = \varepsilon_1 = \varepsilon_2$



Note. Both δ_1, δ_2 parameters are set at $\delta_1 = \delta_2 = 1/T^2$.

decreasing with larger ε_2 values. On the other hand, the impact of ε_1 is not significant or clear. This is expected from the structure of the algorithm, because ε_2 is used in the `PRIVATEMLE` subroutine, which directly affects the model estimates used in subsequent price offerings.

9. Discussion and Insights


9.1. Insufficiency of Input Perturbation

Input perturbation is a straightforward method for designing differentially private algorithms and is actually an effective method in some application scenarios. The high-level idea of input perturbation is to artificially calibrate noise directly to the *inputs* of the

algorithm in order to protect private information. With noisy inputs, the privacy of the entire algorithm trivially follows from the closeness-to-postprocessing property of differential privacy (Fact 2).

In the context of personalized dynamic pricing, the application of the input perturbation method amounts to calibrating noise directly to the personal features x_t of each incoming customer: $\tilde{x}_t = x_t + \omega_t$ for some centered noise vectors $\{\omega_t\}_{t=1}^T$. Such an approach, however, is likely to fail because the features of each individual customer are relatively independent from each other. Therefore, a very large magnitude of noises $\{\omega_t\}$ needs to be injected, which renders the subsequent pricing algorithm impractical. A more detailed discussion follows:

Figure 8. (Color online) Average Regret of Our Proposed Algorithm Under Different δ Parameter Values



Notes. From left to right, the δ values are $1/T, 1/T^2, \dots, 1/T^{10}$. The time horizon is measured in terms of 10^5 periods (i.e., $T = 2$ means 2×10^5 total time periods).

Table 1. Average Regret Comparison with Nonprivate Pricing Algorithms

Algorithm	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.5$	$\varepsilon = 1.0$	$\varepsilon = 5.0$	Nonprivate
$T = 10^5, d = 2$	201×10^{-4}	142×10^{-4}	74.6×10^{-4}	41.9×10^{-4}	44.7×10^{-4}	3.1×10^{-4}
$T = 10^6, d = 2$	46.0×10^{-4}	19.8×10^{-4}	11.5×10^{-4}	5.6×10^{-4}	5.5×10^{-4}	0.6×10^{-4}
$T = 10^5, d = 3$	156×10^{-4}	130×10^{-4}	92.6×10^{-4}	62.9×10^{-4}	43.4×10^{-4}	3.1×10^{-4}
$T = 10^6, d = 3$	74.9×10^{-4}	49.4×10^{-4}	23.1×10^{-4}	14.7×10^{-4}	5.7×10^{-4}	1.6×10^{-4}

1. Suppose $\tilde{x}_t = x_t + w_t$ is the anonymized version of a customer's feature vector x_t at time t . Because \tilde{x}_t is released and used in the subsequent process of the pricing algorithm, one must ensure that \tilde{x}_t is differentially private. This means that the magnitude of w_t must be sufficiently large (on the order of $\Omega(1/\varepsilon)$) to protect the sensitive information of x_t .

2. Usually, input perturbation results in a much worse performance of the differentially private algorithms compared with output perturbation. Consider the very simple example of having sensitive data x_1, \dots, x_n , and one wants to release $\tilde{x} = \frac{1}{n} \sum_{i=1}^n x_i$ with ε -differential privacy. If we use input perturbation with $\tilde{x}_i = x_i + w_i$ and the Laplace mechanism, we have $w_i \sim \text{Lap}(0, 1/\varepsilon)$, and therefore $\tilde{x}^1 := \frac{1}{n} \sum_{i=1}^n \tilde{x}_i$ satisfies $\mathbb{E}[\|\tilde{x}^1 - \tilde{x}\|] \asymp O(1/\varepsilon\sqrt{n})$. On the other hand, if one uses output perturbation by releasing $\tilde{x}^2 := \tilde{x} + \frac{1}{n} \text{Lap}(0, 1/\varepsilon)$, then one has $\mathbb{E}[\|\tilde{x}^2 - \tilde{x}\|] \asymp O(1/\varepsilon n)$. It is easy to verify that both \tilde{x}^1, \tilde{x}^2 are differentially private, but \tilde{x}^2 clearly is much closer to \tilde{x} compared with \tilde{x}^1 . This very simple example shows that, in general, input perturbation (directly adding noises to sensitive data) is usually less efficient and should be avoided if there are better approaches.

3. In the particular model studied in this paper, the use of a generalized linear model further complicates the input perturbation-based methods. For many generalized linear models, such as the logistic regression model, the efficiency of statistical estimates (e.g., the MLE) decays *exponentially* fast with respect to the vector norm of the feature vector x . Hence, if we use $\tilde{x}_t = x_t + w_t$ to replace x_t directly in the logistic regression model, the norm of \tilde{x}_t is on the order of $\Omega(1/\varepsilon)$, and therefore the resulting method is going to incur an $O(\exp\{1/\varepsilon\})$ term in regret, which makes the regret excessively large.

In Table 3 we compare the average regret of our proposed algorithm with the input perturbation

method using numerical simulations. Table 3 shows that the regret of our designed algorithm is significantly smaller than that of the input perturbation. Furthermore, the average regret of input perturbation is very large unless the ε parameter is at least 1 and does not necessarily decrease with increasing number of time periods T .

9.2. Impact of Privacy Constraints on Consumer Surplus

In this section we study the impact of privacy constraints of the seller's personalized pricing algorithm on the average consumer's surplus, under different levels of privacy constraints. We model the utility u_t for each incoming customer at time t with feature vector x_t and offered price p_t as $u_t = \zeta \langle \phi(x_t, p_t), \theta^* \rangle + \zeta_t$, where $\zeta = 4$, $\phi(x_t, p_t) = \frac{1}{\sqrt{d}} [x_t; p_t]$, and ζ_t are i.i.d. random variables following the standard centered logistic distribution. The customer will make one unit of purchase if $u_t > 0$, resulting in a surplus of u_t , and leave without making any purchases if $u_t < 0$, resulting in zero surplus at that period. It is easy to verify that this utility model leads to the logistic regression model we used in the numerical experiments, or more specifically, $\Pr[y_t = 1 | x_t, p_t] = \Pr[u_t > 0 | x_t, p_t] = \frac{e^{\zeta \phi_t^T \theta^*}}{1 + e^{\zeta \phi_t^T \theta^*}}$, where $\phi_t = \phi(x_t, p_t)$.

Figure 9 reports the average consumer surplus under our proposed privacy-aware personalized pricing algorithm, for both $d = 2$ and 3 with consumers' contextual vectors $\{x_t\}_{t=1}^T$ and the unknown regression model synthesized in the same way as in Section 8. We also plot the consumer surplus for a hypothetical pricing algorithm that is *not* subject to any privacy constraints as dashed lines in Figure 9. Note that we did not incorporate consumers' surplus from the protection of their private data, which is difficult to measure and compare against the surplus from their

Table 2. Average Regret with $T = 10^5$ of Our Algorithm Under Different $\varepsilon_1, \varepsilon_2$ Settings

Algorithm	$\varepsilon = 0.02$	$\varepsilon = 0.05$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.5$
$T = 10^5, d = 2$, fix $\varepsilon_1 \equiv 0.1$	0.0247	0.0232	0.0195	0.0145	0.0073
$T = 10^5, d = 2$, fix $\varepsilon_2 \equiv 0.1$	0.0192	0.0178	0.0196	0.0192	0.0191
$T = 10^5, d = 3$, fix $\varepsilon_1 \equiv 0.1$	0.0164	0.0160	0.0147	0.0128	0.0092
$T = 10^5, d = 3$, fix $\varepsilon_2 \equiv 0.1$	0.0145	0.0149	0.0144	0.0149	0.0154

Notes. When the row indicates "fix $\varepsilon_1 \equiv 0.1$ " (or "fix $\varepsilon_2 \equiv 0.1$ "), then the ε in the column represents the value of ε_2 (or accordingly, ε_1).

Table 3. Average Regret of Our Proposed Algorithm and the Input Perturbation Method

Algorithm	Our algorithm			Input perturbation		
	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 1.0$	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 1.0$
$T = 10^5, d = 2 (\times 10^{-4})$	142	74.6	41.9	393	393	98.2
$T = 5 \times 10^5, d = 2 (\times 10^{-4})$	42.7	19.4	10.4	393	393	95.8
$T = 10^6, d = 2 (\times 10^{-4})$	19.8	11.5	5.6	393	393	95.3

purchasing decisions. As we can see from Figure 9, as ϵ increases from 0 to 1, the implied privacy protection becomes *weaker* as the adversary has a stronger ability to distinguish between neighboring databases. This means that as ϵ increases, the seller has less ability to discriminate against customers based on their personal data and features, resembling a transition from first-degree to third-degree price discrimination. As a result, the consumer surplus increases as ϵ increases, and the seller extracts less of the consumer surplus from his or her limited ability to carry out price discrimination.

10. Conclusions and Future Directions

In this paper, we investigate how to protect the privacy of a customer's personal information and purchasing decisions in personalized dynamic pricing with demand learning. Under the generalized linear model of the demand function, we propose a privacy-preserving constrained MLE policy. We establish both the privacy guarantee under the notion of anticipating DP and the regret bounds for oblivious adversarial and stochastic settings.

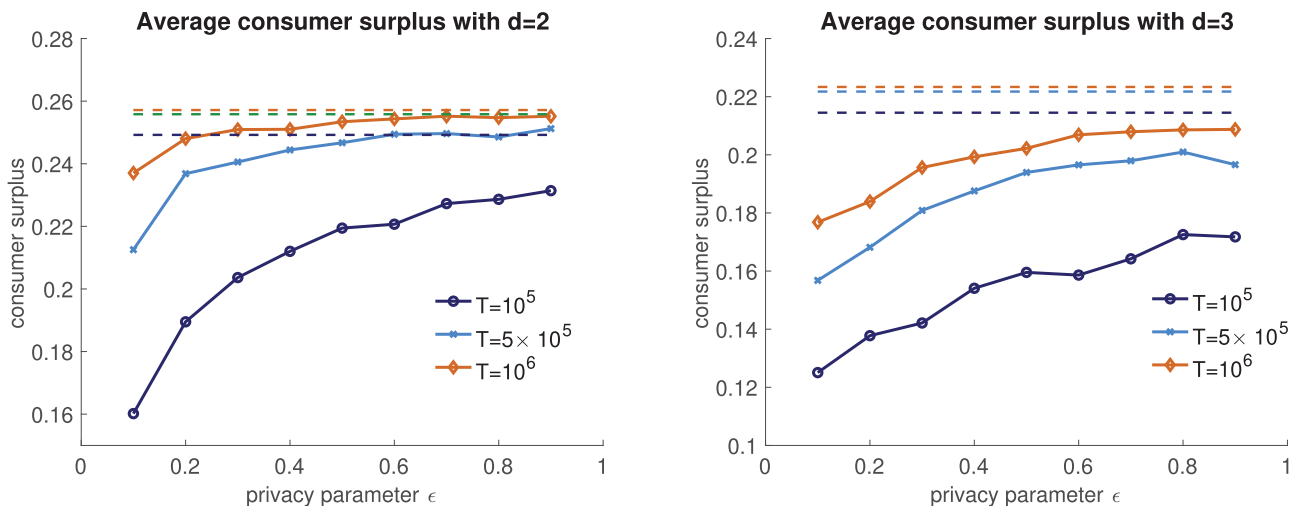
There are several future research directions. First, we could extend the current privacy setting to the

local DP (Evfimievski et al. 2003, Kasiviswanathan et al. 2011), which is a stronger notion of DP. The local DP is suitable for distributed environments, as user terminals need to randomize data before sending it to the center. A very recent paper by Ren et al. (2020) investigates the UCB algorithm under the local DP. It would be interesting to study the personalized dynamic pricing under this stronger notion of DP. More important, as privacy has become a significant concern for the public, especially in the e-commerce domain, we believe that systematic research on privacy-preserving revenue management will become increasingly important in both academia and industry. Although there is relatively less research in this area, we hope our work inspires future studies on privacy-aware operations management (e.g., inventory control or assortment optimization) based on the DP framework.

Acknowledgments

Author names are listed in alphabetical order (Xi Chen is the corresponding author). The authors thank the department editor, the associated editor, and the anonymous referees for many useful suggestions and feedback, which greatly improves the paper.

Figure 9. (Color online) Average Consumer Surplus Under Different Levels of Privacy Constraints and Time Horizons



Notes. The dashed lines represent average consumer surplus for a personalized pricing algorithm *not* subject to any data privacy constraints. Both ϵ_1, ϵ_2 parameters are equal to ϵ in the figures, and both δ_1, δ_2 parameters are set as $1/T^2$, where T is the time horizon.

Endnote

¹ This setting is known as the “oblivious adversary” model in the contextual bandit literature. Although this model is weaker than the “fully adversarial” one mostly studied in the literature, we adopt the oblivious adversary model for a more convenient treatment of privacy constraints, as $\{x_t\}$ will not depend on the offered prices or the randomly realized demands.

References

- Abowd JM (2018) The U.S. Census Bureau adopts differential privacy. *Proc. 24th ACM SIGKDD Internat. Conf. Knowledge Discovery Data Mining* (ACM, New York), 2867.
- Apple Differential Privacy Team (2017) Learning with privacy at scale. Technical report, Apple, Cupertino, CA.
- Araman VF, Caldentey R (2009) Dynamic pricing for nonperishable products with demand learning. *Oper. Res.* 57(5):1169–1188.
- Ban G-Y, Keskin NB (2021) Personalized dynamic pricing with machine learning: High dimensional features and heterogeneous elasticity. *Management Sci.* Forthcoming.
- Barbaro M, Zeller T Jr (2006) A face is exposed for AOL Searcher No. 4417749. *New York Times* (August 9), <https://www.nytimes.com/2006/08/09/technology/09aol.html>.
- Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Oper. Res.* 57(6):1407–1420.
- Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Sci.* 61(4):723–739.
- Broder J, Rusmevichientong P (2012) Dynamic pricing under a general parametric choice model. *Oper. Res.* 60(4):965–980.
- Chan T-HH, Shi E, Song D (2011) Private and continual release of statistics. *ACM Trans. Inform. System Security* 14(3):1–24.
- Chaudhuri K, Monteleoni C, Sarwate AD (2011) Differentially private empirical risk minimization. *J. Machine Learn. Res.* 12(3):1069–1109.
- Chen Q, Jasin S, Duenyas I (2015) Real-time dynamic pricing with minimal and flexible price adjustment. *Management Sci.* 62(8):2437–2455.
- Chen L, Mislove A, Wilson C (2016) An empirical analysis of algorithmic pricing on Amazon Marketplace. *Proc. 25th Internat. Conf. World Wide Web* (International World Wide Web Conferences Steering Committee, Geneva), 1339–1349.
- Chen X, Owen Z, Pixton C, Simchi-Levi D (2021) A statistical learning approach to personalization in revenue management. *Management Sci.*, ePub ahead of print January 18, <https://doi.org/10.1287/mnsc.2020.3772>.
- Cheung WC, Simchi-Levi D (2017) Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. Preprint, submitted November 27, <http://dx.doi.org/10.2139/ssrn.3075658>.
- Cheung WC, Simchi-Levi D, Wang H (2017) Dynamic pricing and demand learning with limited price experimentation. *Oper. Res.* 65(6):1722–1731.
- Cohen MC, Lobel I, Paes Leme R (2020) Feature-based dynamic pricing. *Management Sci.* 66(11):4921–4943.
- Dani V, Hayes TP, Kakade SM (2008) Stochastic linear optimization under bandit feedback. *Proc. 21st Annual Conf. Learn. Theory* (Association for Computational Learning, Sweden), 355–366.
- den Boer AV, Zwart B (2013) Simultaneously learning and optimizing using controlled variance pricing. *Management Sci.* 60(3):770–783.
- Ding B, Kulkarni J, Yekhanin S (2017) Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, vol. 30 (Curran Associates, Red Hook, NY), 3571–3580.
- Dong J, Roth A, Su WJ (2019) Gaussian differential privacy. Preprint, submitted May 7, <https://arxiv.org/abs/1905.02383>.
- Dwork C, Roth A (2014) The algorithmic foundations of differential privacy. *Foundations Trends Theoret. Comput. Sci.* 9(3–4):211–407.
- Dwork C, McSherry F, Nissim K, Smith A (2006a) Calibrating noise to sensitivity in private data analysis. Halevi S, Rabin T, eds. *Theory Cryptography Conf.* (Springer, Berlin), 265–284.
- Dwork C, Naor M, Pitassi T, Rothblum GN (2010) Differential privacy under continual observation. *Proc. 42nd Annual ACM Sympos. Theory Comput.* (ACM, New York), 715–724.
- Dwork C, Talwar K, Thakurta A, Zhang L (2014) Analyze Gauss: Optimal bounds for privacy-preserving principal component analysis. *Proc. 46th Annual ACM Sympos. Theory Comput.* (ACM, New York), 11–20.
- Dwork C, Kenthapadi K, McSherry F, Mironov I, Naor M (2006b) Our data, ourselves: Privacy via distributed noise generation. Vaudenay S, ed. *Adv. Cryptology: Annual Internat. Conf. Theory Appl. Cryptographic Techniques* (Springer, Berlin), 486–503.
- Erlingsson Ú, Pihur V, Korolova A (2014) RAPPOR: Randomized aggregatable privacy-preserving ordinal response. *Proc. 2014 ACM SIGSAC Conf. Comput. Comm. Security* (ACM, New York), 1054–1067.
- Evfimievski A, Gehrke J, Srikant R (2003) Limiting privacy breaches in privacy preserving data mining. *Proc. 22nd ACM SIGMOD-SIGACT-SIGART Sympos. Principles Database Systems* (ACM, New York), 211–222.
- Farias VF, Van Roy B (2010) Dynamic pricing with a prior on market response. *Oper. Res.* 58(1):16–29.
- Ferreira KJ, Simchi-Levi D, Wang H (2018) Online network revenue management using Thompson sampling. *Oper. Res.* 66(6):1586–1602.
- Filippi S, Cappe O, Garivier A, Szepesvári C (2010) Parametric bandits: The generalized linear case. Lafferty J, Williams C, Shawe-Taylor J, Zemel R, Culotta A, eds. *Advances in Neural Information Processing Systems*, vol. 23 (Curran Associates, Red Hook, NY), 586–594.
- Hannak A, Soeller G, Lazer D, Mislove A, Wilson C (2014) Measuring price discrimination and steering on e-commerce web sites. *Proc. 2014 Internet Measurement Conf.* (ACM, New York), 305–318.
- Harrison JM, Keskin NB, Zeevi A (2012) Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Sci.* 58(3):570–586.
- Javanmard A, Nazerzadeh H (2019) Dynamic pricing in high-dimensions. *J. Machine Learn. Res.* 20(9):1–49.
- Kasiviswanathan SP, Lee HK, Nissim K, Raskhodnikova S, Smith A (2011) What can we learn privately? *SIAM J. Comput.* 40(3):793–826.
- Kifer D, Smith A, Thakurta A (2012) Private convex empirical risk minimization and high-dimensional regression. *Proc. 25th Annual Conf. Learn. Theory, JMLR Workshop and Conference Proceedings*, vol. 23 (Association for Computational Learning, Scotland), 25.1–25.40.
- Lei YM, Miao S, Momot R (2020) Privacy-preserving personalized revenue management. HEC Paris Research Paper MOSI-2020-1391, HEC Paris, Jouy en Josas, France.
- Li L, Lu Y, Zhou D (2017) Provably optimal algorithms for generalized linear contextual bandits. *Proc. 34th Internat. Conf. Machine Learn.* (PMLR, Australia), 2071–2080.
- Linden G, Smith B, York J (2003) Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Comput.* 7(1):76–80.
- Miao S, Chen X, Chao X, Liu J, Zhang Y (2019) Context-based dynamic pricing with online clustering. Preprint, submitted November 1, <https://arxiv.org/abs/1902.06199>.
- Mishra N, Thakurta A (2015) (Nearly) optimal differentially private stochastic multi-arm bandits. Meila M, Heskes T, eds. *Proc. 31st*

- Conf. Uncertainty Artificial Intelligence (AUAI Press, Arlington, VA), 592–601.
- Mohammed R (2017) How retailers use personalized prices to test what you're willing to pay. *Harvard Bus. Rev.* (October 20), <https://hbr.org/2017/10/how-retailers-use-personalized-prices-to-test-what-youre-willing-to-pay>.
- Narayanan A, Shmatikov V (2008) Robust de-anonymization of large sparse datasets. *Proc. IEEE Sympos. Security Privacy* (IEEE Computer Society, Washington, DC), 111–125.
- Qiang S, Bayati M (2016) Dynamic pricing with demand covariates. Preprint, submitted June 1, <http://dx.doi.org/10.2139/ssrn.2765257>.
- Ren W, Zhou X, Liu J, Shroff NB (2020) Multi-armed bandits with local differential privacy. Preprint, submitted July 6, <https://arxiv.org/abs/2007.03121>.
- Shariff R, Sheffet O (2018) Differentially private contextual linear bandits. Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, eds. *Advances in Neural Information Processing Systems*, vol. 31 (Curran Associates, Red Hook, NY), 4296–4306.
- Tang W, Ho C-J, Liu Y (2020) Differentially private contextual dynamic pricing. *Proc. 19th Internat. Conf. Autonomous Agents MultiAgent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC), 1368–1376.
- Tringale M (2018) Dynamic pricing vs. personalized pricing, what's the difference? *Wiser* (blog), July 5, <https://blog.wiser.com/dynamic-pricing-vs-personalized-pricing-whats-the-difference/>.
- Tsitsiklis J, Xu K, Xu Z (2021) Private sequential learning. *Oper. Res.*, ePub ahead of print January 14, <https://doi.org/10.1287/opre.2020.2021>.
- Wang Y-X (2019) Per-instance differential privacy. *J. Privacy Confidentiality* 9(1), <https://doi.org/10.29012/jpc.662>.
- Wang Z, Deng S, Ye Y (2014) Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Oper. Res.* 62(2):318–331.
- Wang Y-X, Lei J, Fienberg SE (2016) On-average KL-privacy and its equivalence to generalization for max-entropy mechanisms. Domingo-Ferrer J, Pejić-Bach M, eds. *Proc. Internat. Conf. Privacy Statist. Databases* (Springer, Cham, Switzerland), 121–134.
- Wang Y, Chen X, Chang X, Ge D (2020) Uncertainty quantification for demand prediction in contextual dynamic pricing. *Production Oper. Management*, ePub ahead of print December 9, <https://doi.org/10.1111/poms.13337>.
- Wang Y, Wang R, Du SS, Krishnamurthy A (2019) Optimism in reinforcement learning with generalized linear function approximation. Preprint, submitted December 9, <https://arxiv.org/abs/1912.04136>.
- Xu K (2018) Query complexity of Bayesian private learning. Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R, eds. *Advances in Neural Information Processing Systems*, vol. 31 (Curran Associates, Red Hook, NY), 2431–2440.
- Xu J, Xu K, Yang D (2020) Optimal query complexity for private sequential learning against eavesdropping. Preprint, submitted August 17, <https://arxiv.org/abs/1909.09836>.
- Zheng K, Cai T, Huang W, Li Z, Wang L (2020) Locally differentially private (contextual) bandits learning. Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H, eds. *Advances in Neural Information Processing Systems*, vol. 33 (Curran Associates, Red Hook, NY), 12300–12310.