

Structured Dynamic Pricing: Optimal Regret in a Global Shrinkage Model

Rashmi Ranjan Bhuyan

BHUYANR@USC.EDU

*Department of Data Sciences and Operations, Marshall School of Business
University of Southern California, Los Angeles, CA 90089, USA*

Adel Javanmard

AJAVANMA@USC.EDU

*Department of Data Sciences and Operations, Marshall School of Business
University of Southern California, Los Angeles, CA 90089, USA*

Sungchul Kim

SUKIM@ADOBE.COM

Adobe Research, 345 Park Avenue, San Jose, CA 95110 USA

Gourab Mukherjee

GOURAB@USC.EDU

*Department of Data Sciences and Operations, Marshall School of Business
University of Southern California, Los Angeles, CA 90089, USA*

Ryan A. Rossi

RYROSSI@ADOBE.COM

Adobe Research, 345 Park Avenue, San Jose, CA 95110 USA

Tong Yu

TYU@ADOBE.COM

Adobe Research, 345 Park Avenue, San Jose, CA 95110 USA

Handong Zhao

HAZHAO@ADOBE.COM

Adobe Research, 345 Park Avenue, San Jose, CA 95110 USA

Editor: Qiang Liu

Abstract

We consider dynamic pricing strategies in a streamed longitudinal data set-up where the objective is to maximize, over time, the cumulative profit across a large number of customer segments. We consider a dynamic model with the consumers' preferences as well as price sensitivity varying over time. Building on the well-known finding that consumers sharing similar characteristics act in similar ways, we consider a global shrinkage structure, which assumes that the consumers' preferences across the different segments can be well approximated by a spatial autoregressive (SAR) model. In such a streamed longitudinal set-up, we measure the performance of a dynamic pricing policy via regret, which is the expected revenue loss compared to a clairvoyant that knows the sequence of model parameters in advance. We propose a pricing policy based on penalized stochastic gradient descent (PSGD) and explicitly characterize its regret as functions of time, the temporal variability in the model parameters as well as the strength of the auto-correlation network structure spanning the varied customer segments. Our regret analysis results not only demonstrate asymptotic optimality of the proposed policy but also show that for policy planning it is essential to incorporate available structural information as policies based on unshrunk models are highly sub-optimal in the aforementioned set-up. We conduct simulation experiments across a wide range of regimes as well as real-world networks based studies and report encouraging performance for our proposed method.

Keywords: Dynamic pricing, Regret, Stochastic Gradient Descent, Spatial Models, Log-likelihood

1. Introduction

Due to the ubiquitous reach of digital marketing, dynamic pricing settings are extensively studied by firms that sell a significant fraction of their inventories over online marketplaces and through digital advertisements (see Cohen et al., 2020; Keskin and Zeevi, 2016; Ban and Keskin, 2021; Bimpikis et al., 2019; Javanmard, 2017; Leme and Schneider, 2018; Javanmard and Nazerzadeh, 2019 and the references therein). As such it is a vibrant topic of research in online machine-learning (Zhou et al., 2019; Cesa-Bianchi et al., 2015), operations research (Golrezaei et al., 2020; Cheung et al., 2017), information (Cui et al., 2021), marketing (Schwartz et al., 2017; Choi et al., 2020) and management sciences (Farias and Van Roy, 2010; Broder and Rusmevichientong, 2012; den Boer and Zwart, 2013). For fuller references see Sec 1.3.

In this work, we study the problem of a firm selling a product to customers who arrive over time. The firm has the opportunity to set different prices not only over time t but also for different customer segments $l = 1, \dots, L$. We consider setting the prices across these customer segments in a dynamic manner such that the expected cumulative revenue, aggregated over the customer segments as well as time is maximized. As a motivational example, consider the digital marketing problem (Liu-Thompkins, 2019) where an advertisement of a product priced at p_{lt} is shown to n_{lt} customers in segment l at time t . Let $y_{ltk}, k = 1, \dots, n_{lt}$ denote the binary variables corresponding to conversion based on the advertisement, i.e., $y_{ltk} = 1$ if the k -th advertisement in the l -th segment at time t led to a purchase, and $y_{ltk} = 0$ otherwise. Often in these problems, the firm also has the opportunity to access other covariates $\mathbf{x}_{lt}$ such as demographic information for the customer segment l at time t .

As time t progresses, the goal is to explore and set prices p_{lt} optimally based on the current covariates $\mathbf{x}_{lt}$ as well as on the previous customer responses $\{y_{lsk} : 1 \leq s < t\}$ and their associated prices and covariate information. The goal is to optimize the cumulative revenue

$$\sum_{t=1}^T \sum_{l=1}^L y_{lt} p_{lt} \text{ where, } y_{lt} = \sum_{k=1}^{n_{lt}} y_{ltk}. \quad (1)$$

1.1 Streamed Longitudinal Probit Set-up

Demand heterogeneity (Bimpikis et al., 2019; Chintagunta et al., 2002) is traditionally tackled by segmenting consumers who have similar purchasing propensity as well as similar responses to price changes. Though truly homogeneous segments of consumers do not exist, the approximation provides a reasonable interface to design differential pricing strategies that optimally target each customer segments. Modern online trading platforms, marketplaces and lead generation systems, facilitate implementing price differential strategies across a wide range of segments. Often advertisers have access to the geographical location of the consumer and these segments based on zip-codes of the consumers (Train, 2009). Another popular choice is segmenting customers based on the different marketing channels by which they were

approached (Berman and Thelen, 2018). In accordance with these modern applications, we consider the number of segments L to be large. In the existing literature (Javanmard, 2017), the overall revenue in (1) is optimized for probability models based on rational choice theory, which assumes that consumers are rational and make choices that maximize their utility.

Let the utility function U_{ltk} for the k th customer in the l th segment at time t be given by the following additive model:

$$U_{ltk} = \alpha_{lt} + \beta_t p_{lt} + \mathbf{x}'_{lt} \boldsymbol{\mu}_t + \sigma Z_{ltk} \quad (2)$$

where $k = 1, \dots, n_{lt}$; α s are the preferences of the customers that vary across both time and segments; β s are the price-sensitivities of the customer. Vectors $\boldsymbol{\mu}$ are coefficients corresponding to the non-priced covariates and vary over time but invariant across segments. Z s are independent and identically from Gaussian distribution with mean 0 and has variance 1. Based on rational choice theory, if $U_{ltk} > 0$ then a sale occurs, i.e., $Y_{ltk} = 1$; else, $Y_{ltk} = 0$. Let $Y_{lt} = \sum_{k=1}^{n_{lt}} Y_{ltk}$ be the count of sales for segment l at time t . Then,

$$Y_{lt} \sim \text{Binomial}(n_{lt}, q_{lt}),$$

where $q_{lt} = \Phi(\sigma^{-1}(\alpha_{lt} + \beta_t p_{lt} + \mathbf{x}'_{lt} \boldsymbol{\mu}_t))$ and Φ is the cumulative distribution function of standard Gaussian distribution. Note, that q_{lt} is not only a function of the model parameters but also depends on the price. We use capital letters for random variables, small letters for the values a random variable takes, and boldface letters for vectors and matrices.

Note that at each time t we advertise the product at price p_{lt} based on the covariates $\mathbf{x}_{lt}$ and observe the total sales across the segment Y_{lt} . It is crucial to note that we do not directly observe the utilities associated with each segment. This diverges from the classical bandit framework, where complete losses, or utilities in this case, are observed. Instead, in our scenario, we only receive partial information, specifically whether the utility was positive or negative depending on whether a sale was made. Based on the sales information, the joint log-likelihood across all segments at time t is given by

$$\ell_t(\boldsymbol{\lambda}) = \sum_{l=1}^L y_{lt} \log q_{lt} + (n_{lt} - y_{lt}) \log(1 - q_{lt}), \quad (3)$$

where $\boldsymbol{\lambda} = \{\boldsymbol{\lambda}_t := (\boldsymbol{\alpha}_t, \beta_t, \boldsymbol{\mu}_t) : t = 1, \dots, T\}$ and $\boldsymbol{\alpha}_t = (\alpha_{1t}, \dots, \alpha_{Lt})$. The expected revenue from segment l subjected to price p_{lt} at time t is,

$$\begin{aligned} \text{Rev}(\boldsymbol{\lambda}, l, t, p_{lt}) &= p_{lt} \mathbb{E}_{\boldsymbol{\lambda}}(Y_{lt}) = n_{lt} p_{lt} q_{lt} \\ &= n_{lt} p_{lt} \Phi(\sigma^{-1}(\boldsymbol{\alpha}_{lt} + \beta_t p_{lt} + \mathbf{x}'_{lt} \boldsymbol{\mu}_t)), \end{aligned}$$

and the goal is to maximize the cumulative revenue

$$\text{Rev}(\boldsymbol{\lambda}, \mathbf{p}) = \sum_{l=1}^L \sum_{t=1}^T \text{Rev}(\boldsymbol{\lambda}, l, t, p_{lt})$$

over the prices $\mathbf{p} = \{p_{lt} : 1 \leq l \leq L, 1 \leq t \leq T\}$ that the firm can set. Conditioned on the parameters $\boldsymbol{\lambda}$, maximizing $\text{Rev}(\boldsymbol{\lambda}, \mathbf{p})$ decouples into separate maximization of the revenue of

each segment at each time point. The first-order condition for optimal price p_{lt}^* conditional on the set of parameters is

$$p_{lt}^* = -\sigma \beta_t^{-1} \frac{\Phi(\sigma^{-1}(\alpha_{lt} + \beta_t p_{lt}^* + \mathbf{x}_{lt}' \boldsymbol{\mu}_t))}{\phi(\sigma^{-1}(\alpha_{lt} + \beta_t p_{lt}^* + \mathbf{x}_{lt}' \boldsymbol{\mu}_t))}. \quad (4)$$

As p_{lt}^* depends on the unknown model parameters $\boldsymbol{\lambda}_t$, we call this the oracle price and $\text{Rev}(\boldsymbol{\lambda}, \mathbf{p}^*)$ imposes the highest theoretically achievable upper bound on the revenue. For any other pricing policy $\mathbf{p}$ we define its regret over the oracle strategy as:

$$\begin{aligned} \mathcal{R}(\boldsymbol{\lambda}, \mathbf{p}) &= \sum_{l=1}^L \sum_{t=1}^T \mathcal{R}_{lt}(\boldsymbol{\lambda}, \mathbf{p}), \text{ where} \\ \mathcal{R}_{lt}(\boldsymbol{\lambda}, \mathbf{p}) &= \text{Rev}(\boldsymbol{\lambda}, l, t, p_{lt}^*) - \text{Rev}(\boldsymbol{\lambda}, l, t, \hat{p}_{lt}). \end{aligned}$$

1.2 SAR based global Shrinkage Structures

We impose the following regularity condition on the temporal changes in the price sensitivity and the covariate effects:

$$\sum_{t=1}^{T-1} |\beta_{t+1} - \beta_t| \leq C_\beta^* \text{ and } \sum_{t=1}^{T-1} \|\boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t\|_2 \leq C_\mu^*. \quad (5)$$

Note that these cumulative temporal changes in price sensitivity and covariate effects C_β^* and C_μ^* are not constant and can grow over time. Unlike the price sensitivity, the preference coefficients $\boldsymbol{\alpha}_t$ however greatly depends on the state of the current inventory and can highly fluctuate over time. However, it is well known that people who are close to each other in some networks often reflect highly corrected preferences (Bradlow et al., 2005; Ma et al., 2015). Spatial models provide a natural way to model this correlation between different units of analysis based on their contiguity in a network (Banerjee et al., 2014; Gelfand et al., 2010; LeSage, 2004). Geographic closeness is a proxy for many socio-demographic variables like income, education, wealth and property values, which are also related to consumer purchase behavior, and has been the primary focus of a large number of existing pricing models (Yang and Allenby, 2003; Jank and Kannan, 2005; Bimpikis et al., 2019). Networks based on non-geographic metrics can also capture preference similarities among customers (Karmakar et al., 2021). Consider the following *Spatially Autoregressive* (SAR) structure (see ch. 6 of Anselin, 2013 and ch. 2 of Banerjee et al., 2014) on the α_{lt} :

$$\alpha_{lt} = \rho_t \sum_{j=1}^L w_{lj} \alpha_{jt} + \tau \epsilon_{lt}, \quad (6)$$

where $w_{lj} \geq 0$, ϵ_{lt} are i.i.d $N(0, 1)$ and $\tau > 0$. In its most basic form, (6) imposes a global hierarchical structure with the auto-correlation parameter ρ_t regulating the level of global spillovers (and hence connectedness) among the units. Relation (6) implies having the following hierarchical prior on the preference parameters:

$$\boldsymbol{\alpha}_t \sim N_L(\mathbf{0}, \tau^2 (\mathbf{I} - \rho_t \mathbf{W})^{-2}). \quad (7)$$

We consider the network structure and its associated contiguity matrix $\mathbf{W}$ to be invariant over time. SAR models such as above have been very successful in assimilating spatial network information in real-world datasets (Manski, 1993; Anselin, 2013). In Bramoullé et al. (2009), SAR is used in modeling recreational services consumption by secondary school students, whereas Hsieh and Lee (2016) used SAR to incorporate friendship networks of high school students in predicting their academic performances. In Zhou et al. (2017), SAR is used to model user activity on social media regarding transportation services in China. We allow the auto-correlation parameter ρ_t to vary over time while satisfying the regularity condition

$$\sum_{t=1}^{T-1} |\rho_{t+1} - \rho_t| \leq C_\rho^*. \quad (8)$$

Since $\boldsymbol{\alpha}_t$ is a random variable, our defined regret $\mathcal{R}(\boldsymbol{\lambda}, \mathbf{p})$ is also random. We define a parameter set $\boldsymbol{\theta} = \{\boldsymbol{\theta}_t := (\rho_t, \beta_t, \boldsymbol{\mu}_t) : t = 1, \dots, T\}$ which we utilize to establish a Bayes regret, based on the prior on $\boldsymbol{\alpha}_t$ in equation (7) as follows:

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) = \mathbb{E}\{\mathcal{R}(\boldsymbol{\lambda}, \mathbf{p})\}, \quad (9)$$

where the expectation is over randomness associated with the distribution of α_t governed by (7). Let Θ be a set of parameter $\boldsymbol{\theta}$ satisfying (5) and (8). For this set of parameters, we consider developing dynamic pricing strategies $\mathbf{p}$ that minimize the Bayes regret in (9).

1.3 Our Contributions and Related Work

We develop a Projected Stochastic Gradient Descent (PSGD) algorithm based on the logarithm of the marginal likelihood $\ell_t(\boldsymbol{\theta}) = \log\{\mathbb{E}\{\exp \ell_t(\boldsymbol{\lambda})\}\}$ which is the convolution of the likelihood in (3) with the prior in (7). We show that the proposed algorithm controls the Bayes regret at the order of $\mathcal{O}(\sqrt{T})$. We also show that for any data-driven pricing strategy the Bayes regret can not be of the lower order of $\mathcal{O}(\sqrt{T})$. Thus, as $T \rightarrow \infty$, the proposed algorithm is asymptotically rate-optimal. Our main result, Theorem 2 is provided in Section 3.

An important attribute of Theorem 2 is that, we provide an explicit characterization of the Bayes regret of the proposed PSGD algorithm in terms of not only time T but also as functions of the model parameters and the underlying heterogeneity (difference in the n_{lt}) in the data. We show how the regret of the proposed algorithm depends on temporal variability in the model parameters as well as on the strength of correlation among the segments. Our upper-bound on the regret of the prescribed method (see (23)) depends on the spectral radius of the SAR structure in (9). It is sensitive to the magnitude of the autocorrelation parameter and greatly contracts as the correlation increases.

In Corollary 6, we show that any unshrunk pricing policy that does not borrow strength across the customer segments is highly sub-optimal with respect to the proposed strategy. This is in accordance with classical statistical shrinkage theory results (Fourdrinier et al., 2018) that are based on non-dynamic set-ups. To see the connections consider the penalized likelihood criterion:




Figure 1: Schematic on the working principle of proposed PSGD. The red path are the PSGD updates on the marginal likelihood with penalty.

$$\text{PL}(\lambda; \omega) = \sum_{t=1}^T \{ \ell_t(\lambda) + \omega \|(I - \rho_t W)\alpha_t\|_2^2 \}. \quad (10)$$

Running a vanilla stochastic gradient descent (with projection on Θ) based on this penalized criteria is asymptotically equivalent to applying the proposed PSGD algorithm on the marginal log-likelihood. However, the same algorithm based on the unpenalized likelihood $\text{PL}(\lambda; 0)$ will have higher estimation error in the estimates of α_t when L is large, which would in turn yield a significantly higher regret. In this context, it is crucial for any decent pricing policy to shrink its α_t estimates towards the ellipsoids $\{\alpha_t : \|(I - \rho_t W)\alpha_t\|_2 \leq s_\omega\}$. Figure 1 shows the schematic for this essential shrinkage effect on the α_t . The rigorous mathematical proof is provided in Corollary 6.

Our research is connected to and builds on recent works in statistical shrinkage theory, online machine learning and econometrics theory on demand modeling. Next, we list the relevant literature in these research and also briefly mention our contributions.

1) Dynamic Pricing with Online Learning There exists a growing body of research on dynamic pricing with learning (den Boer, 2015; Farias and Van Roy, 2010; Harrison et al., 2012; Cesa-Bianchi et al., 2015; Ferreira et al., 2018; Cheung et al., 2017). The classical formulations of this problem Broder and Rusmevichientong (2012); den Boer and Zwart (2013); Besbes and Zeevi (2009) consider parametric model for the demand-price curve, which is unknown and the learner aims to learn, via exploration-exploitation of prices, while aiming to obtain a low regret in revenue. These works focus on non-contextual settings (no features for customers), and are relevant to applications where a seller is offering an unlimited supply of a single product to the market. Recently, there was significant interest in contextual-models, which use the customers and products attributes to model willingness-to-pay of the buyers for the products, potentially in a heterogeneous way and offer personalized pricing (Leme and Schneider, 2018; Cohen et al., 2020; Ban and Keskin, 2021; Javanmard and Nazerzadeh, 2019; Lobel et al., 2016; Golrezaei et al., 2021). In addition, some of the recent works in this area (Javanmard, 2017; Keskin and Zeevi, 2017) aim to model the temporal behavior of buyers, by considering time-dependent demand models.

Closer to our analysis is the notion of dynamic regret, which has been used in online convex optimization to evaluate the performance of a learner against a dynamic target Zinkevich (2003); Yang et al. (2016); Jadbabaie et al. (2015); Besbes et al. (2015). However, the general framework of online optimization does not directly apply to our setting, since in the former framework, after taking an action each step the learner observes the incurred loss (or some first order information on it), which can be used in next rounds. In contrast, in our setting the firm only observes the count of sales at each segment, and not the utility of customers. Our work is the first to propose and analyze a dynamic contextual demand model, which accounts for temporal behavior of consumers as well as the network effect among them via spatially autoregressive structure.

2) Dynamic Hierarchical modeling. Hierarchical modeling provides an effective tool for pooling information across similar units and is one of the most popular approaches for modeling large and complex data sets (Fourdrinier et al., 2018; Banerjee et al., 2014; Kou and Yang, 2017). Here, (7) imposes a dynamic hierarchical structure on the customer preference coefficients that are linked through a time-invariant non-exchangeable network structure in the second-level prior on (2). Recent applications of hierarchical modeling to consumer responses in digital products Banerjee et al. (2022); Mukhopadhyay et al. (2022); Banerjee et al. (2021) have been very successful in analyzing structured longitudinal data-sets abet in a non-dynamic set-up. Here, we provide an extensive characterization on the operational characteristics of PSGD in a streamed longitudinal set-up and thereby provide theoretical support for the popular PSGD approach for hierarchical modeling in dynamic set-ups.

3) Modeling Demand Heterogeneity. Spatial models are very popular in operations management and information sciences to capture non-stationarity in demand (Karmakar et al., 2021; Bimpikis et al., 2019; Jank and Kannan, 2005). In (2)-(7) we have a dynamic spatial model that is governed by changes in the auto-correlation parameter. An important feature of our contributions is that we show PSGD is able to track the variation of the auto-correlation parameter over time and yield asymptotically rate-optimal regret in a dynamic spatial model.

4) Shrinkage prediction under heterogeneity. It is now commonplace to use notions of shrinkage to improve predictive performances of algorithms in multi-parametric set-ups (Hastie et al., 2009; Efron and Hastie, 2021). Recent results in (Xie et al., 2012; Tan, 2015; Weinstein et al., 2018; Brown et al., 2018) have brought to light new shrinkage phenomena in heteroscedastic models. Here, we study shrinkage prediction in a heteroscedastic dynamic set-up; n_{lt} – the number of customers in segment l approached at time t can greatly vary over l and t . This is an important aspect of our model. It entertains high imbalance across the design matrix but greatly increases applicability. Particularly, in “pull” marketing systems unlike “push” systems (Peter et al., 2000) the firm has no control on the number of customers who visits the site/store and does a price check. Thus, n_{lt} will be large in some zip-codes/demographics and quite low in others. Due to the SAR structure in (7) it is possible to learn the preferences α_{lt} with high precision even in segments with very low n_{lts} . In Section 3 and 4, we illustrate the impact of the heterogeneous n_{lts} on the regret bounds.

5) Spatial models and applications. Spatial models provide a natural way to model the correlation between different units of analysis based on how close they are in a similarity

space. Spatial models based on correlated customer preferences have been successfully employed in marketing and economics to model real-world sales data with high predictive accuracy. Jank and Kannan (2005) used a spatially correlated preference model to model consumer choices of two product forms of a book –print or PDF. Yang and Allenby (2003) estimate a binary choice model akin to ours in which consumer preferences for a vehicle’s country-of-origin (Japanese/non-Japanese) are spatially correlated based on the distance and demographic similarity between consumers. Ma et al. (2015) models consumers’ decision of whether or not to purchase a callback ringtone. In several non-marketing data applications also, spatial autoregressive (SAR) models have been very successful in assimilating spatial network information in real-world datasets (Manski, 1993; Anselin, 2013). Bramoullé et al. (2009) used SAR to model consumption of recreational services such as participation in artistic, sports and social activities by secondary school students, whereas Hsieh and Lee (2016) used SAR to incorporate the friendship networks of high school students to predict academic performance. Zhou et al. (2017) used SAR to model user activity on social media regarding transportation services in China. Based on this existing literature which shows that SAR can well capture the correlation among customers and users in economic and social real-world data, we feel that the proposed model will be good for real-world applications in dynamic set-up.

2. Proposed PSGD Algorithm

2.1 Assumptions

We make some assumptions on the covariate and the parameter space to simplify the presentation of our results. The covariates are normalized such that $\|\mathbf{x}_{it}\| \leq 1$. Similarly the parameters $\boldsymbol{\mu}_t$ are such that $\|\boldsymbol{\mu}_t\| \leq C_\mu$ where C_μ is a known constant. This gives a ball of radius C_μ in which the parameters reside. We can even allow the parameter to belong in any convex set Θ_μ . The results would then depend on the size of the parameter space up to a constant factor.

Based on (2), we also assume that the price sensitivity β_t should be negative i.e. an increment in price decreases the utility of the product for the consumer. We also make an assumption on the lower and the upper bound on the magnitude of price sensitivity $c_\beta \leq |\beta_t| \leq C_\beta$. These restrictions inherently create the restricted space Θ_μ and Θ_β for our model parameters.

We also make two key assumptions on the SAR structure and auto-correlation parameter.

Assumption 2.1. $\mathbf{W}$ is a symmetric, positive semidefinite (PSD) kernel e.g. radial basis function (RBF) kernel.

Since $\mathbf{W}$ denotes a distance matrix across the L segments, it is natural that $\mathbf{W}$ is a symmetric matrix. It is also worth noting that the common choices of kernels in non-parametric estimation are PSD (see Tsybakov (2008), Section 1.2).

Assumption 2.2. The interaction parameter ρ_t for all time periods is positive and uniformly bounded away from the reciprocal of the maximum eigenvalue of the interaction matrix i.e. $\exists \varepsilon > 0$, such that $\rho_t \leq (1 - \varepsilon)/\omega^*$, where ω^* is the largest eigenvalue of the known interaction matrix $\mathbf{W}$.

The spatial autoregressive structure of the preference coefficients α_t in equation (6) with the Gaussian noises implies the joint Gaussian nature of the preference coefficient with variance $(\mathbf{I} - \rho_t \mathbf{W})^{-2}$ in (7). Since covariance matrices are always positive semi definite, so $\mathbf{I} - \rho_t \mathbf{W}$ is positive, hence, $\rho_t \leq 1/\omega^*$. With this assumption, we imply the inequality to be strict. Otherwise, the model becomes degenerate (i.e. covariance of Gaussian distribution becomes rank-deficient) and in that case one can work with the lower-dimensional space where the SAR covariance is full-rank.

2.2 Reparameterization

The hierarchical prior in (7) can be used to write the explicit value of α_{lt} in terms of prior hyperparameters ρ_t , τ and standard multivariate normal noise ϵ as $\alpha_{lt} = \tau \langle \mathbf{e}_l, (\mathbf{I} - \rho_t \mathbf{W})^{-1} \epsilon \rangle$, where $\mathbf{e}_l$ denotes the l^{th} basis vector. Substituting α_{lt} in the utility model, we can rewrite (2) in terms of $\boldsymbol{\theta}$ as

$$U_{ltk} = \beta_t p_{lt} + \mathbf{x}'_{lt} \boldsymbol{\mu}_t + \sigma Z_{ltk} + \tau \langle \mathbf{e}_l, (\mathbf{I} - \rho_t \mathbf{W})^{-1} \epsilon \rangle. \quad (11)$$

This produces a marginal model, where the utility for each segment l can be described as a normal distribution with variance $V_{lt}^2 = \|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 \tau^2 + \sigma^2$. Next, normalizing the utility to have unit variance we consider the following reparameterized utility model

$$\tilde{U}_{ltk} = b_{lt} p_{lt} + \mathbf{x}'_{lt} \mathbf{m}_{lt} + Z_{ltk}, \quad (12)$$

where, $b_{lt} = \beta_t/V_{lt}$ and $\mathbf{m}_{lt} = \boldsymbol{\mu}_t/V_{lt}$. We use this marginal utility model for describing our policy.

2.3 Optimal Pricing

Since the noises in (4) are distributed as standard Gaussian, it follows that the optimal price p_{lt}^* is the solution to the equation:

$$p_{lt}^* = -b_{lt}^{-1} \frac{\Phi(b_{lt} p_{lt}^* + \mathbf{x}'_{lt} \mathbf{m}_{lt})}{\phi(b_{lt} p_{lt}^* + \mathbf{x}'_{lt} \mathbf{m}_{lt})}. \quad (13)$$

The optimality condition in (13) can be restructured as $\varphi(-b_{lt} p_{lt}^* - \mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt} = 0$ where $\varphi(v) = v - \Phi(-v)/\phi(v)$ is the virtual valuation function (Myerson, 1981). With the use of the valuation function, we can explicitly describe the optimal price p_{lt}^* as a function of the utility model parameters

$$p_{lt}^* := g(b_{lt}, \mathbf{m}_{lt}) = -\frac{\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}}{b_{lt}}. \quad (14)$$

Proposition 1 Consider the definition of marginal variance $V_{lt}^2 = \|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 \tau^2 + \sigma^2$. Under Assumptions 2.1 and 2.2, the variance V_{lt}^2 satisfies

$$c_V \leq V_{lt} \leq C_V,$$

where $c_V^2 = \tau^2 + \sigma^2$ and $C_V^2 = \tau^2/\varepsilon^2 + \sigma^2$. Additionally, the optimal prices satisfy $p_{lt}^* \leq M$, where $M = c_\beta^{-1} C_V (C_\mu c_V^{-1} - 0.5\phi(0))$.

Algorithm 1 PSGD based Dynamic Pricing Policy**Data** $\mathbf{W}$: known segment structure**Initialize** $p_{l1} = c$, $\hat{b}_{l1} \in \Theta_b$ and $\hat{\mathbf{m}}_{l1} \in \Theta_m \forall l$ **for** $t = 1, 2, \dots$ **do** **Data** $y_{lt}, \mathbf{x}_{l,t+1}$: Longitudinal data stream 1. Compute the gradient $\mathcal{L}'_{lt}(\hat{\mathbf{m}}_{lt}, \hat{b}_{lt})$ (16) for each segment l 2. Update parameters by moving in the opposite direction of gradient with step size η_t and then projecting onto the restricted space (19)

$$\hat{b}_{l,t+1} = \Pi_{\Theta_b}(\hat{b}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^b); \quad \hat{\mathbf{m}}_{l,t+1} = \Pi_{\Theta_m}(\hat{\mathbf{m}}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^m)$$

 3. Set price $p_{l,t+1}$ using the optimal pricing function $p_{l,t+1} = g(\hat{b}_{l,t+1}, \hat{\mathbf{m}}_{l,t+1})$ **2.4 Proposed Pricing Policy**

We propose a pricing policy based on a projected stochastic gradient descent on the loss function described in (15). With the PSGD, we aim to estimate the reparameterized parameter set $(b_{lt}, \mathbf{m}_{lt})$ for every segment.

Based on the assumptions $\|\boldsymbol{\mu}_t\| \leq C_\mu$ and $|\beta_t| \leq C_\beta$ we first define $\Theta_b := \{\beta/c_V : \beta \in \Theta_\beta\}$ and $\Theta_m := \{\boldsymbol{\mu}/c_V : \boldsymbol{\mu} \in \Theta_\mu\}$, the restricted space of the new parameters $b, \mathbf{m}$. These are natural extensions to the assumptions since $b_{lt} = \beta_t/V_{lt}$ and c_V is the lower bound on V_{lt} .

Next, define the loss function as the negative of the log-likelihood function for the utility model (12).

$$\mathcal{L}_t(\boldsymbol{\lambda}) = - \sum_{l=1}^L y_{lt} \log q_{lt} + (n_{lt} - y_{lt}) \log(1 - q_{lt}), \quad (15)$$

with $q_{lt} = \Phi(b_{lt} p_{lt} + \mathbf{x}'_{lt} \mathbf{m}_{lt})$. Each summand in the loss (15) is the loss for a specific segment l . Let $\mathcal{L}_{lt}$ denote these losses,

$$\mathcal{L}_{lt} = -y_{lt} \log q_{lt} - (n_{lt} - y_{lt}) \log(1 - q_{lt}).$$

We compute the gradient of loss functions: $\mathcal{L}'_{lt}(b_{lt}, \mathbf{m}_{lt}) = (\nabla \mathcal{L}_{lt}^{(b)}, \nabla \mathcal{L}_{lt}^{(m)})$ for each segment as

$$\nabla \mathcal{L}_{lt}^{(b)} = -y_{lt} \frac{\phi(u_{lt}^0)}{\Phi(u_{lt}^0)} + (n_{lt} - y_{lt}) \frac{\phi(-u_{lt}^0)}{\Phi(-u_{lt}^0)} p_{lt}, \quad (16)$$

$$\nabla \mathcal{L}_{lt}^{(m)} = -y_{lt} \frac{\phi(u_{lt}^0)}{\Phi(u_{lt}^0)} + (n_{lt} - y_{lt}) \frac{\phi(-u_{lt}^0)}{\Phi(-u_{lt}^0)} \mathbf{x}_{lt}. \quad (17)$$

At time point t , we move in the opposite direction of the gradient with step size η_t . The resultant estimates are then projected onto the restricted space Θ_b, Θ_m based on the assumptions on the size of the parameters to get the successive estimates:

$$\hat{b}_{l,t+1} = \Pi_{\Theta_b}(\hat{b}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^b), \quad (18)$$

$$\hat{\mathbf{m}}_{l,t+1} = \Pi_{\Theta_m}(\hat{\mathbf{m}}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^m), \quad (19)$$

where, $\Pi_{\Theta_b}(\cdot)$ and $\Pi_{\Theta_m}(\cdot)$ are the projection functions on to the convex set Θ_b and Θ_m respectively. The policy finally uses these estimated parameters and the optimal pricing function $g(\cdot, \cdot)$ defined in (14) to set the price for the next period. The method is summarized in Algorithm 1.

Remark: Our valuation model assumes the knowledge of interaction matrix $\mathbf{W}$, yet in many real-world applications, access to such matrices is not readily available. In these scenarios, we can utilize covariates that describe the segments to construct the interaction matrix. However, please note that, these covariates however can not be used in the utility model (2) again as that would make the covariates in the utility model and $\mathbf{W}$ correlated which is not covered in the paper. An example of this is demonstrated in our simulations, Section 5 (setups 3-7), where the segments consist of 48 US states, and we use 15 demographic and socio-economic variables such as age, average income, and unemployment rates to create the interaction matrix $\mathbf{W}$.

Once we possess the covariates that define the segments, we can employ kernel functions on these covariates to define a similarity matrix $\mathbf{W}$. These kernel techniques require estimation of a bandwidth (tuning parameter), which we can determine using cross-validation. To perform cross-validation, we utilize past sales data where the prices were not dynamically changed and create a train/test split. Next we select a fixed bandwidth to generate an interaction matrix $\tilde{\mathbf{W}}$, and estimate our model on the training data based on $\tilde{\mathbf{W}}$. Finally, based on the log-likelihood of our fitted model on the test split, we select the optimal bandwidth.

3. Theoretical Results

In this section, we provide the bounds on the regret for the dynamic pricing policy we employ. We show that under regularity conditions on the temporal nature of the parameters β , μ and ρ , the regret as defined in (9) has order square root of the time horizon T . We also show the optimality of the bound by showing that no policy can achieve a worst-case regret better than the same rate.

3.1 Upper Bound on Regret of the Proposed Algorithm

We present an upper bound on the regret of the proposed pricing policy in terms of the reparameterized model parameters in (12). Later, through lemma 3, we create a link between the actual parameters and the reparameterized ones. Finally, we show that when the step sizes $\eta_t \propto 1/\sqrt{t}$, we achieve $\mathcal{O}(\sqrt{T})$ regret.

Theorem 2 *For any $\theta \in \Theta$ defined below (9), the regret of our proposed policy satisfies:*

$$\mathcal{B}(\theta, p) \leq \mathcal{R}_1 + \mathcal{R}_2 + \mathcal{R}_3 + \mathcal{R}_4 + \mathcal{O}(\log T), \text{ where,} \quad (20)$$

$$\begin{aligned}
\mathcal{R}_1 &= C_1 \sum_{t=1}^T \sum_{l=1}^L \eta_t^{-1} |b_{l,t+1} - b_{lt}|, \\
\mathcal{R}_2 &= C_2 \sum_{t=1}^T \sum_{l=1}^L \eta_t^{-1} \|\mathbf{m}_{l,t+1} - \mathbf{m}_{l,t}\|_2, \\
\mathcal{R}_3 &= C_3 \sum_{t=1}^T \sum_{l=1}^L \eta_t n_{lt}^2 \leq C_3 \sum_{t=1}^T \eta_t n_t^2, \text{ and,} \\
\mathcal{R}_4 &= C_4 \eta_{T+1}^{-1} L.
\end{aligned}$$

C_1, C_2, C_3 and C_4 are constants independent of T, n, L and the model parameters.

The detailed proof of the theorem is presented in the Appendix. We present a brief overview of its proof in Section 4. We next concentrate on further explaining the terms on the right side of (20). We simplify the first two terms $\mathcal{R}_1, \mathcal{R}_2$ in Theorem 2 and provide a key lemma 3. For ease of notation we define $\delta_{t\nu} = \|\nu_{t+1} - \nu_t\|$ for any parameter ν . This helps us transform our regret from the reparameterized quantities $b, \mathbf{m}$ to the original parameters $\beta, \boldsymbol{\mu}$ and ρ .

Lemma 3 *Let ω_* be the smallest eigenvalue of $\mathbf{W}$, then under Assumptions 2.1 and 2.2, the variation across the parameters in the utility model (12), can be bounded as*

$$|b_{l,t+1} - b_{lt}| \leq \tau^{-1}(1 - \rho_t \omega_*) \delta_{t\beta} + C_5 \delta_{t\rho}, \quad (21)$$

$$\|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 \leq \tau^{-1}(1 - \rho_t \omega_*) \delta_{t\mu} + C_5 \delta_{t\rho}. \quad (22)$$

Next, we demonstrate the implications of the above result in a simplified setup with any network structure $\mathbf{W}$. The goal is to understand the effect of the network structure and the auto-correlation (ρ_t) on the upper bound of the regret in (20). We provide the following corollary that explicitly shows the relation of regret with the auto-correlation parameter.

Corollary 4 *If $\eta_t \propto 1/\sqrt{t}$, and $\rho_t = \rho$ for all t , then the dynamic pricing policy based on Algorithm 1 has regret*

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) \leq C_6 \tau^{-1} (1 - \rho \omega_*) \sum_{t=1}^T \sqrt{t} (\delta_{t\beta} + \delta_{t\mu}) + \mathcal{O}(\sqrt{T}). \quad (23)$$

Corollary 4 shows that the regret has two parts, one with order $\sqrt{T}$, while the other part depends on the temporal nature of price sensitivity and customer preferences. The regret occurred in this part depends on the strength of the network inversely, i.e, higher the strength of the network (higher the ρ) lower the regret and vice-versa.

Note that if ρ_t was varying across time, we can extend the bound in Corollary 4. Assume that $\rho_* = \min_t \rho_t$, then the bound on regret can be modified as

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) \leq C_6 \tau^{-1} (1 - \rho_* \omega_*) \sum_{t=1}^T \sqrt{t} (\delta_{t\beta} + \delta_{t\mu}) + C_7 \sum_{t=1}^T \sqrt{t} \delta_{t\rho} + \mathcal{O}(\sqrt{T}), \quad (24)$$

where C_6 and C_7 are constants. The bound above behaves similarly to Corollary 4 if the temporal changes across auto-correlation is small.

3.2 Lower Bound on Regret of Any Data-driven Policy

We show that the bound in corollary 4 is indeed tight in terms of dependence on the time horizon. In the next theorem we show that there exists parameters in the space Θ such that under the demand model with these parameters, the regret of any policy is of the order at least $\sqrt{T}$. The detailed proof is provided in the Appendix.

Theorem 5 *Consider the utility model (12) and let $N_T := \sum_{t=1}^T \sum_{l=1}^L n_{lt}$ be the total number of costumers across all segment and times up to T . For any fixed graph $\mathbf{W}$, the worst-case risk of any data driven pricing policy $\hat{\mathbf{p}}$ satisfies*

$$\min_{\hat{\mathbf{p}}} \max_{\boldsymbol{\theta} \in \Theta} \mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) \geq C_8 \sqrt{T} (1 + \log(N_T/T)),$$

for some constant C_8 . In particular, if $n_{lt} \geq 1$ for all l, t , we have

$$\min_{\hat{\mathbf{p}}} \max_{\boldsymbol{\theta} \in \Theta} \mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) \geq C_8 \sqrt{T} (1 + \log(L)).$$

Theorem 5 along with (24) implies that our pricing policy in algorithm 1, is optimal, if the temporal changes across the price sensitivity, customer preferences and auto-correlation is of the order $\sqrt{T}$, i.e. if $\sum_{t=1}^T \sqrt{t}(\delta_{t\beta} + \delta_{t\mu} + \delta_{t\rho}) = \mathcal{O}(\sqrt{T})$, then our policy is order optimal. Below we give examples of two such scenarios where our pricing policy is order optimal. The first example illustrates a scenario with smooth variations, and the second one illustrates a scenario with bursty changes.

Consider a case where the core parameters $\boldsymbol{\mu}, \beta$ and ρ are perturbed at each time t such that $\delta_{t\mu} \propto 1/t$, $\delta_{t\beta} \propto 1/t$ and $\delta_{t\rho} \propto 1/t$. In such a scenario, the upper bound on regret using (24) is of the order $\sum_{t=1}^T \sqrt{t}(\delta_{t\beta} + \delta_{t\mu} + \delta_{t\rho}) + \mathcal{O}(\sqrt{T})$. Since the perturbations across all the parameters is inversely proportional to t , the sum $\sum_{t=1}^T \sqrt{t}(\delta_{t\beta} + \delta_{t\mu} + \delta_{t\rho})$ simplifies to $\sum_{t=1}^T 1/\sqrt{t}$ which is $\mathcal{O}(\sqrt{T})$. In this case, both the upper bound and lower bound are of $\mathcal{O}(\sqrt{T})$ showing optimality. Please note that, in this setting we can replace the order of perturbations to any function of t smaller than $1/t$ and the regret of our policy will still be optimal.

Next, consider a case where $\delta_{t\mu} = c$, $\delta_{t\beta} = c$ and $\delta_{t\rho} = c$ for some constant c whenever $t \in \{2^0, 2^1, 2^2, \dots\}$ and 0 otherwise. In this discrete case, the size of perturbations is the same while the duration across the perturbations keeps increasing. Again, using (24), the upper bound on regret is of the order $\sum_{t=1}^T \sqrt{t}(\delta_{t\beta} + \delta_{t\mu} + \delta_{t\rho}) + \mathcal{O}(\sqrt{T})$. The sum $\sum_{t=1}^T \sqrt{t}(\delta_{t\beta} + \delta_{t\mu} + \delta_{t\rho})$ can be bounded by $c \int_0^{\log_2 T} 2^{x/2} dx$ which is $\mathcal{O}(\sqrt{T})$. Thus, our pricing policy is optimal in this discrete case as well.

3.3 Sub-optimality of Unshrunk Pricing Policies

Next, we consider unshrunk policies that do not incorporate the structure (7) on the $\boldsymbol{\alpha}_t$ s. Such unshrunk policies suffer from severe noise accumulation in estimating $\boldsymbol{\alpha}_t$ as free parameters at every time point. The following result whose proof is provided in Section C.4 of the appendix shows that the Bayes regret from any unshrunk pricing policies based on the unpenalized likelihood is highly sub-optimal as compared to the proposed strategy $\mathbf{p}$. Consider a parametric space $\bar{\Theta}$ such that any $\boldsymbol{\theta} \in \bar{\Theta}$ satisfies that $\sum_{t=1}^T \sqrt{t}\delta_{t\beta}$, $\sum_{t=1}^T \sqrt{t}\delta_{t\mu}$

and $\sum_{t=1}^T \sqrt{t} \delta_{t\rho}$ are $\mathcal{O}(\sqrt{T})$. The following result shows sub-optimality of unshrunk pricing policies over $\bar{\Theta}$.

Lemma 6 *For any $\theta \in \bar{\Theta}$, the regret of any data-driven policy $\mathbf{p}_U$ based on the unpenalized likelihood in (10) satisfies*

$$\mathcal{B}(\theta, \mathbf{p}_U) / \mathcal{B}(\theta, \mathbf{p}) = \Omega(\sqrt{T}).$$

4. Outline of Proofs and Overview of Techniques

For the detailed proofs of all the results, we refer to the appendix. In this section, we delve into the intuition and the intermediate steps used in proving the two main results in Section 3.

4.1 Proof Sketch of Theorem 2

The crucial idea is to bound the revenue loss (regret) with the parameters in the model. To achieve that consider the revenue function with the utility model (12)

$$\text{Rev}_{lt}(p_{lt}) = n_{lt} p_{lt} \Phi(b_{lt} p_{lt} + \mathbf{x}'_{lt} \mathbf{m}_{lt}). \quad (25)$$

The revenue loss using our policy is then the difference between $\text{Rev}_{lt}(p_{lt}^*)$ and $\text{Rev}_{lt}(p_{lt})$ where p_{lt}^* is the optimal price for the model true parameters and p_{lt} is the price posted with the dynamic pricing policy.

Proposition 7 *There exists a constant C_9 such that the regret of our policy on segment l at time t can be bounded as*

$$\mathcal{R}_{lt} = \text{Rev}_{lt}(p_{lt}^*) - \text{Rev}_{lt}(p_{lt}) \leq C_9 n_{lt} (p_{lt} - p_{lt}^*)^2, \quad (26)$$

where p_{lt} is our posted price and p_{lt}^* is the optimal price that maximizes the revenue under known parameters.

We simplify the regret bound term $(p_{lt} - p_{lt}^*)^2$ on the right hand side of (26) in our next lemma. The idea is to use the optimal pricing function $g(\cdot, \cdot)$ defined in section 2.3. The prices p_{lt}^* and p_{lt} can then be defined as $p_{lt}^* = g(b_{lt}, \mathbf{m}_{lt})$, the optimal price based on the true parameters and $p_{lt} = g(\hat{b}_{lt}, \hat{\mathbf{m}}_{lt})$, the optimal price with respect to the estimated parameters that our proposed policy posts. The lemma then hinges on the fact that the function $g(\cdot, \cdot)$ defined in (14) is Lipschitz.

Lemma 8 *For model (12), under the true parameters $b_{lt}, \mathbf{m}_{lt}$ and the output $\hat{b}_{lt}, \hat{\mathbf{m}}_{lt}$ from our PSGD pricing policy, the following holds true:*

$$(p_{lt} - p_{lt}^*)^2 \leq C_{10} \langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + C_{10} p_{lt}^2 (b_{lt} - \hat{b}_{lt})^2 \quad (27)$$

for some constant $C_{10} > 0$.

The $\mathcal{R}_{lt}$ terms in (26) are the building blocks for our total regret $\mathcal{B}(\theta, \mathbf{p})$ as in (9) where $\mathcal{B}(\theta, \mathbf{p}) = \sum_{t=1}^T \sum_{l=1}^L \mathcal{R}_{lt}$. The above two lemmas relate the revenue regret occurred by the policy with the estimation error of the parameters in the model (12). The final step involves creating a link between this estimation error and the temporal nature of the parameters to achieve the regret bound as in Theorem 2.

4.2 Proof Sketch of Theorem 5

For the lower bound, we want to find worst case scenarios in terms of parameters. In this case, we use the idea of “uninformative prices” (Broder and Rusmevichientong, 2012). These are prices where the purchase probability curves for all different sets of parameters intersect. Such prices do not reveal any information about the parameters since all the purchase curves contain the point. These uninformative prices become an issue when they are also the optimal prices for some set of parameters. If a policy wants to learn the parameters fast, they need to do exploration away from these uninformative prices. But during the process of exploration it chooses parameters farther from the actual parameters and thus increases regret.

The general idea of such proofs is to create a setting where these optimal prices are indeed uninformative. In the proof, we show existence of such parameters and their corresponding optimal “uninformative prices”. Calling these parameters γ_0 , the proof hinges on two relations, one showing that learning the utility model parameters closely is expensive in terms of regret:

$$\text{Reg}_T^{\pi, \gamma_0} \geq \frac{C_{12}}{(\gamma_0 - \gamma)^2} \text{KL}(f_T^{\pi, \gamma_0}; f_T^{\pi, \gamma}) \quad (28)$$

where $f_t^{\pi, \gamma}$ is the density of purchases for all consumers until time t , provided that the policy π is employed. An interpretation of the KL-divergence $\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma})$ is the certainty level of the policy π about the true model parameters γ_0 , over some other counterfactual parameter γ . So the above bound implies that increasing certainty about the underlying model is costly.

The next bound shows that if the policy can not differentiate between two parameters that are “close” to each other (in other words to increase its confidence in one), then again it incurs a large regret. Specifically, if $\gamma_1 = \gamma_0 + 1/(4T^{1/4})$,

$$\text{Reg}_T^{\pi, \gamma_0} + \text{Reg}_T^{\pi, \gamma_1} \geq C_{13} \sqrt{T} e^{-\text{KL}(f_T^{\pi, \gamma_0}; f_T^{\pi, \gamma_1})}. \quad (29)$$

Intuitively, the first equation shows that exploitation is necessary (choosing the optimal parameter γ_0 to have small KL-divergence) and the second one asks for exploration to stay away from uninformative prices, so as to gain information about the model parameters and increase the certainty about it, as measured by KL-divergence.

5. Numerical Experiments

We study the performance of the proposed algorithm using numerical experiments based on synthetic as well as real-world based networks. We consider a wide range of regimes with varying (a) temporal variation of the model parameters (b) strength and nature of the network (c) sampling heterogeneity across sectors, and (d) noise distributions.

Set-up 1. Consider $L = 10$ segments and a time-invariant sampling policy with different sampling rates across two segment groups: for any $t \geq 1$, $n_{lt} = 50$ for $l = 1, \dots, 5$ and $n_{lt} = 200$ for $l = 6, \dots, 10$. We use a time-invariant network $\mathbf{W}$ that was generated using radial basis function (RBF) kernel of width one on independent standard Gaussian feature vectors, drawn from input space $\mathbb{R}^{10}$. We use bivariate covariates $\mathbf{x}_{lt}$ generated from standard exponential distribution and set $\tau = 1, \sigma = 1$ in (2)-(6). The price sensitivity and the customer preferences are assumed as $\beta_1 = -0.4$ and $\boldsymbol{\mu}_1 = (0.1, 0.15)$. With change in time the parameters change as follows, $\beta_{t+1} = \beta_t + \delta_{t\beta}$; $\boldsymbol{\mu}_{t+1} = \boldsymbol{\mu}_t + \boldsymbol{\delta}_{t\mu}$, where $\delta_{t\beta} = t^{-b} \tilde{Z}_t / (10|\tilde{Z}_t|)$




Figure 2: Regret plots of the proposed policy in Set-up 1 as b which governs the shift in the model parameters over time changes. The plots for $b = 0.5, 1, \infty$ are in red, green and blue respectively. In panels (a) Regret in original scale (b) $\log(\text{Regret})$ vs $\log(T)$. The dotted lines in panel (b) are the best fitted line with slope 0.5.

and $\delta_{t\mu} = t^{-b} \tilde{\mathbf{Z}}_t / (10 \|\tilde{\mathbf{Z}}_t\|)$ where $\tilde{\mathbf{Z}}_t$ and $\bar{\mathbf{Z}}_t$ are standard Gaussian random variables of dimension 1 and 2 respectively.

We set ρ_t to 0.5 for all $t \geq 1$ and consider three cases, $b = 0.5, 1, \infty$, for the temporal variations across β_t and μ_t . Note that the case of $b = \infty$ corresponds to the scenario where the parameters do not change over time.

Recall from Corollary 4 that the regret is upper bounded as

$$\mathcal{B}(\theta, \mathbf{p}) \leq C_6 \tau^{-1} (1 - \rho \omega_*) \sum_{t=1}^T \sqrt{t} (\delta_{t\beta} + \delta_{t\mu}) + \mathcal{O}(\sqrt{T}).$$

For $b = \infty$, the first two terms equate to 0 and the regret is thus $\mathcal{O}(\sqrt{T})$. In the second case when $b = 1$, then $\delta_{t\mu} \propto 1/t$, $\delta_{t\beta} \propto 1/t$. The second term in the sum is thus $\mathcal{O}(\sum_{t=1}^T 1/\sqrt{t})$ which is $\mathcal{O}(\sqrt{T})$ as well. Finally, when we set $b = 0.5$, $\delta_{t\mu} \propto 1/\sqrt{t}$, $\delta_{t\beta} \propto 1/\sqrt{t}$ and the second term is $\mathcal{O}(\sum_{t=1}^T \sqrt{t}/\sqrt{t}) = \mathcal{O}(T)$.

In Figure 2, we plot the regret (cumulative revenue lost to the oracle policy) over time for the three cases. From the figures it is evident that when $b = 0.5$, the regret from the proposed method eventually grows linearly whereas in the other two cases it is controlled at $\mathcal{O}(\sqrt{T})$.

Set-up 2. Here, we aim to study the performance of the proposed method as the strength of the network varies due to change in auto-correlation parameter. We consider a simpler setting than set-up 1 with $L = 4$ segments with a homogeneous sampling rate $n_{lt} = 50$ for all l, t . We consider $b = 1, \rho_t = \rho$ for all $t \geq 1$ and vary $\rho = 0.1, 0.3$ and 0.5 across 3 experiments. Note that, the first two terms $\mathcal{R}_1$ and $\mathcal{R}_2$ in Theorem 2 depend on ρ while the third term $\mathcal{R}_3$ increases with n_{lt} . This means that as n_{lt} increase, $\mathcal{R}_3$ grows very large and the effect of $\mathcal{R}_1$ and $\mathcal{R}_2$ (effect of ρ) on the regret is significantly less. Hence, to see the effect of ρ_t , we consider moderate n_{lt} here. We plot the regret in Figure 3 (a) and see a significant improvement in terms of regret as ρ increases gradually. In Figure 3 (b) we fix $\rho = 0.5$ and compare the regret of our policy with the unshrunk policy based on (10). We see that the regret of the unshrunk policy has linear trend and is quite sub-optimal as T exceeds 20000.

Set-up 3. In this set-up, unlike set-up 1 and 2, we do not consider a randomly generated synthetic network but consider a real-network based on US census data Bureau (2008). Here,




Figure 3: Regret plots for Set-up 2 (a) Regret of proposed policy for $\rho = 0.5$ (blue), 0.3 (green), 0.1 (red) varies (b) Regret of proposed policy (continuous line) and an unshrunk policy (dotted line) for $\rho = 0.5$.

we consider $L = 48$ segments. Each segments constitute a US state. For ease of analysis and presentation, we remove Hawaii and Alaska from the analysis. We choose 15 demographic and socio-economic variables such as percentage of residents in the age group 5-65, average income, unemployment rates, etc for making the network matrix $\mathbf{W}$ among the L segments. We use an RBF kernel of width two and threshold the resultant network at 0.05 level, i.e., edges with weight less than 0.05 are deleted from the network. Figure 4 shows the network.

We generate the covariates $\mathbf{x}_{lt}$ using standard exponential distribution. We apply model (1)–(2) in the main paper in the context of conversion from leads. Consider the problem where at time t , the firm purchases leads n_{lt} for segment l from other lead generation companies and shows pricing p_{lt} which lead to conversion Y_{lt} . We consider that the number of leads n_{lt} is fixed over time t . Let $n_t = n$. For this set-up, we vary $n = 1000, 2500, 5000, 10000, 20000$. We set

$$n_{lt} \propto \text{Population size}_l \times \text{Median Income}_l,$$

i.e., the number of leads in each state is proportional to the population of the state as well as the median disposable income per household in the state. Similar to the analysis in Section 5, we consider the three cases, $b = 0.5, 1, \infty$, for the temporal variations across β_t and μ_t . In Figure 5, we plot the regret for the three cases. We see that the results perfectly match the results in Set-up 1 where we had a random network: with $b = 0.5$, the regret grows linearly, whereas with $b = 1$ and $b = \infty$, the regret is $\mathcal{O}(\sqrt{T})$.

Set-up 4. In the previous setting we assumed that the number of leads in each segment is proportional to the population and the income levels. But for most firms it is highly unlikely that they have good penetration and market share in all the US states. We consider the scenario where the firm has a stronger customer base in some states compared to others. Thus, there will be difference in leads across states. Particularly, in states l where the firm has low penetration, n_{lt} s will be very low. An interesting attribute of this exercise is that the presence of the network structure in (1)–(8) makes the preferences correlated and the preference coefficients from states with low n_{lt} s can also be efficiently learnt by the prescribed method by leveraging the information from states with high penetration.

To study this through numerical experiments, we create an imbalanced design. We divide the states into two groups L_1 and L_2 of equal sizes. Here, n_{lt} are not only proportional to population and income as in Set-up 3 but states in L_1 are given more weightage than those

Table 1: Performance of the prescribed method relative to Unshrunk policy in Set-up 4 as $v_1(1 - v_1)$ varies across columns. (negative implies worse performance)

n	T	0.7	0.8	0.9
1000	100	-1.3%	-4.4%	-0.8%
	500	3.2%	-0.8%	-0.8%
	1000	15.0%	13.9%	12.5%
	5000	50.5%	50.5%	48.5%
2500	100	-2.9%	-5.0%	-5.3%
	500	2.2%	0.8%	-1.4%
	1000	17.1%	14.8%	11.8%
	5000	51.2%	50.3%	48.2%
5000	100	-2.2%	-5.4%	-3.1%
	500	5.1%	1.7%	-0.1%
	1000	18.0%	15.0%	12.0%
	5000	52.3%	49.8%	48.1%
10000	100	-2.4%	-4.9%	-3.8%
	500	3.0%	1.1%	-1.4%
	1000	16.3%	15.0%	10.8%
	5000	51.3%	50.3%	47.5%
20000	100	-1.9%	-4.9%	-4.4%
	500	4.5%	1.1%	-1.5%
	1000	17.3%	15.0%	10.9%
	5000	51.8%	50.3%	47.7%



Figure 4: A network on $L = 48$ US states barring Hawaii and Alaska. This network is based on similarity between states across 15 demographic and economic variables and is thresholded at 0.05. The network is used in experiments for set-ups 3 and 4.

in L_2 , i.e., $v_1 = \sum_{l \in L_1} n_{lt} (\sum_{l \in L_2} n_{lt})^{-1} > 1$. With the same network structure as in set-up 3, we compare the regret of the proposed policy compared to an unshrunk policy at three different levels of v_1 . The first level is $0.7 : 0.3$. The second and third level of imbalance is $0.8 : 0.2$ and $0.9 : 0.1$.

We compute the cumulative regret for all three imbalanced cases for different values of total number of customers $n_t = n$ and time horizon T . We report the relative regret (in terms of percentage) by our proposed policy against the canonical unshrunk policy in Table 1. Relative to the unshrunk policy, the performance of our PSGD based policy is always observed to be superior for moderately large T . The relative performance at higher imbalances is still very high.




Figure 5: Regret of the proposed policy in Set-up 3 for different values of b .

We also report the performance of our policy in these imbalanced design compared to an unshrunk policy in a balanced design regime. We compute the regret of the unshrunk policy in the balanced setting as n_t varies and compare with the efficacy of the prescribed PSGD based policy in the imbalanced setting in Table 2. We see that with this imbalanced design our policy still outperforms the unshrunk policy in a balanced design regime.

Set-up 5. Next, we want to study the regret behaviour under a different network on the segments. Unlike Set-ups 3 and 4, we create a network on the US states based only on the demographic variables. Based on the demographic variables, we create the similarity matrix $\mathbf{W}$ using an RBF kernel of width two and threshold the edges at 0.05. Figure 6 shows the network.

We study the regret of the imbalanced design in this network regime. Similar to the set-up 4, we compare our policy in the imbalanced setting against (a) unshrunk policy in the balanced setup and (b) unshrunk policy in the imbalanced setting. The results are presented in Table 3. At smaller T 's our policy performs worse than an unshrunk policy, but as T grows larger, our policy performs significantly better (more than 50%) compared to the unshrunk policy, both with balanced and imbalanced design.

Set-up 6. We use a network that is based on the similarity across economic variables only. We create the network of the US states based only on the economic variables using an RBF kernel of width two and thresholding the edges at 0.05. Figure 7 shows the network.

We study the performance under this network regime in the imbalanced setting with the two scenarios as above. The results are reported in Table 4. We see that overall, in all imbalanced settings our policy performs much better than the unshrunk policy in the imbalanced as well as the balanced setting.

Set-up 7. With the $\mathbf{W}$ used in Set-ups 3 and 4, we setup an extremely unbalanced design with total number of customers fixed at $n_t = 1000$. In this setting, 10 states have 5 leads each, while the remaining 950 leads are distributed similarly among the remaining 38 states. We specifically study the regret from the 10 states with very low leads.

Consider the two cases here (a) the ten states with very-low-leads are chosen such that they are least connected states (sum of the edge weights is least) in the network (b) the ten states with very-low-leads are chosen such that they are the most connected 10 states

Table 2: Performance of the prescribed method in imbalanced designs relative to Unshrunk policy in balanced design in Set-up 4 as $v_1(1 - v_1)$ varies across columns. (negative implies worse performance)

n	T	0.7	0.8	0.9
1000	100	-6.6%	-7.6%	-16.3%
	500	-1.9%	-9.2%	-22.0%
	1000	10.5%	6.1%	-6.4%
	5000	47.9%	44.9%	35.8%
2500	100	-8.3%	-18.2%	-35.9%
	500	-2.9%	-12.8%	-32.1%
	1000	12.8%	4.1%	-13.6%
	5000	48.7%	43.4%	31.8%
5000	100	-7.6%	-20.1%	-40.1%
	500	0.1%	-19.1%	-36.9%
	1000	13.7%	-2.4%	-19.1%
	5000	49.8%	38.4%	28.8%
10000	100	-7.8%	-128.6%	-36.1%
	500	-2.1%	-127.6%	-33.7%
	1000	11.9%	-95.8%	-17.2%
	5000	48.8%	-16.8%	29.9%
20000	100	-7.2%	-16.5%	-39.3%
	500	-0.5%	-16.6%	-36.4%
	1000	13.0%	-0.1%	-19.0%
	5000	49.3%	40.3%	29.1%

Table 3: Performance of PSGD in Set-up 5 relative to (a) unshrunk policy in the balanced designs and (b) unshrunk policy in the imbalanced designs (negative implies worse performance)

n	T	0.7		0.8		0.9	
		Balanced	Imbalanced	Balanced	Imbalanced	Balanced	Imbalanced
1000	100	-9.1%	-12.4%	-8.2%	-21.7%	-6.1%	-34.0%
	500	-13.5%	-16.9%	-18.0%	-35.5%	-20.8%	-59.1%
	1000	-1.0%	-4.1%	-3.2%	-20.5%	-4.4%	-42.7%
	5000	54.8%	53.4%	54.3%	50.0%	55.7%	49.8%
2500	100	-5.4%	-8.6%	-2.1%	-21.1%	3.2%	-32.5%
	500	-13.8%	-17.2%	-10.2%	-27.4%	-10.8%	-53.6%
	1000	-3.9%	-7.0%	0.5%	-13.8%	0.9%	-40.0%
	5000	45.1%	43.4%	48.7%	49.9%	51.5%	48.8%
5000	100	-8.0%	-11.2%	-7.3%	-19.3%	-2.3%	-31.3%
	500	-15.6%	-19.0%	-15.9%	-27.9%	-13.2%	-45.5%
	1000	-3.8%	-6.9%	-3.3%	-13.9%	0.7%	-30.3%
	5000	46.8%	45.2%	48.3%	49.7%	52.2%	52.3%
10000	100	-7.9%	-11.2%	-4.3%	-22.1%	0.6%	-30.0%
	500	-12.4%	-15.8%	-12.0%	-30.1%	-10.4%	-47.9%
	1000	-0.4%	-3.4%	1.3%	-15.7%	3.3%	-34.5%
	5000	48.5%	47.0%	51.0%	50.6%	54.0%	51.5%
20000	100	-6.9%	-10.1%	-2.4%	-21.0%	1.8%	-31.3%
	500	-13.9%	-17.3%	-10.4%	-29.6%	-10.8%	-49.6%
	1000	-2.1%	-5.2%	1.9%	-16.0%	2.2%	-36.2%
	5000	46.2%	44.6%	49.9%	48.8%	52.5%	50.6%

Table 4: Performance of prescribed method in Set-up 6 compared to (a) unshrunk policy in the balanced setup and (b) unshrunk policy in the imbalanced setup (negative implies worse performance)

n	T	0.7		0.8		0.9	
		Balanced	Imbalanced	Balanced	Imbalanced	Balanced	Imbalanced
1000	100	-4.2%	-2.1%	-7.3%	-2.3%	-4.3%	-3.8%
	500	4.9%	6.8%	2.0%	6.6%	3.6%	4.3%
	1000	10.7%	12.5%	8.8%	12.6%	10.3%	10.3%
	5000	29.3%	30.7%	27.7%	30.8%	29.1%	29.2%
2500	100	-6.2%	-4.0%	-6.2%	-3.0%	-9.4%	-5.9%
	500	4.8%	6.7%	3.4%	6.9%	-1.1%	3.0%
	1000	11.7%	13.5%	10.1%	12.9%	6.4%	9.4%
	5000	30.8%	32.2%	28.9%	31.2%	25.9%	28.5%
5000	100	-7.4%	-5.2%	-6.1%	-2.5%	-9.8%	-5.6%
	500	3.7%	5.6%	3.4%	6.9%	-1.0%	3.0%
	1000	10.2%	12.0%	10.2%	13.0%	6.2%	9.5%
	5000	29.9%	31.3%	29.1%	31.4%	25.9%	28.7%
10000	100	-6.5%	-4.4%	-2.4%	-2.8%	-4.7%	-6.1%
	500	4.5%	6.4%	7.0%	6.7%	3.7%	2.9%
	1000	11.0%	12.8%	13.4%	12.8%	10.8%	9.4%
	5000	30.0%	31.4%	31.4%	31.0%	29.8%	28.9%
20000	100	-5.8%	-3.7%	-3.5%	-2.8%	-6.8%	-5.7%
	500	4.6%	6.5%	6.1%	6.9%	1.6%	3.0%
	1000	11.0%	12.8%	12.7%	13.1%	8.7%	9.3%
	5000	30.2%	31.6%	31.0%	31.4%	28.0%	28.6%





Figure 6: A network on 48 US states (barring Hawaii and Alaska) based on demographic variables. It is used in set-up 5.

in the network. In Figure 8, we plot the relative regret of the prescribed policy from the very-low-leads states with respect to an unshrunk pricing policy.

From the figure we see that in the case (b), the prescribed policy pools information from the other connected states and perform much better compared to an unshrunk policy. On the other hand in case (a) as the low-lead-states are not very well connected with the other states, the prescribed policy perform on a similar level compared to the unshrunk policy and does not provide any benefit.

Set-up 8. Throughout our experiments so far, we have assumed that the noises are Gaussian. Here, we study the performance of our proposed algorithm when the noise in model (2) follows Laplace distribution.

We consider the same synthetic data regime as in setup-1. We have $L = 10$ segments where for any $t \geq 1$, $n_{lt} = 50$ for $l = 1, \dots, 5$ and $n_{lt} = 200$ for $l = 6, \dots, 10$. The network




Figure 7: A network on 48 US states (barring Hawaii and Alaska) based on economic variables. It is used in set-up 6.

matrix $\mathbf{W}$ is generated using an RBF kernel of width one on independent standard Gaussian feature vectors, drawn from input space $\mathbb{R}^{10}$. The covariates $\mathbf{x}_{lt}$ are generated from standard exponential distribution and set v_t to 0.5 for all $t \geq 1$. The price sensitivity and the customer preferences are assumed as $\beta_1 = -0.4$ and $\boldsymbol{\mu}_1 = (0.1, 0.15)$. We simply generate Z_{lt} in (2) as standard Laplace observations rather than Gaussian observations.

We study effects of the temporal variations across β_t and $\boldsymbol{\mu}_t$ by varying b as before. Recall that the case of $b = \infty$ corresponds to the scenario where the parameters do not change over time. Since the noise distribution is Laplace (log-concave), following from Corollary 4, the regret for the two cases of $b = 1, \infty$ should be of order $O(\sqrt{T})$, while the regret for $b = 0.5$ should be $O(T)$. In Figure 9, we plot the regret (cumulative revenue lost to the oracle policy) over time for the three cases. From the figures it is evident that when $b = 0.5$, is linear where as in the other two cases its is controlled at $O(\sqrt{T})$.




Figure 8: Relative regret of the prescribed policy (red) compared to unshrunk policy (green) in Set-up 7. Panel (a) corresponds to case (a) where the very low lead states are not very well connected with other states. Panel (b) corresponds to case (b) where the very low lead states are very well connected with other states.




Figure 9: Regret plots of the proposed policy in Set-up 8 as b which governs the shift in the model parameters over time changes. The plots for $b = 0.5, 1, \infty$ are in red, green and blue respectively. In panels (a) Regret in original scale (b) $\log(\text{Regret})$ vs $\log(T)$. The dotted lines in panel (b) are the best fitted line with slope 0.5.

Set-up 9. We extend our study to a setup where the noises are i.i.d. from Student's t distribution. The setup of this experiment is same as set-up 8, with the only change being the distribution of the noise terms. Here, noises follow Student's t distribution in place of the Laplace distribution. We study the regret of our proposed policy under this setting by varying b . We see the effect of the temporal variations of parameters in Figure 10. We see similar growth of regret, where for $b = 1, \infty$, the regret is $\mathcal{O}(\sqrt{T})$ and for $b = 0.5$ the regret grows linearly. While theoretically the results hold for Gaussian noises, these experiments show that our policy is applicable to a broader family of noise distributions.

6. Discussion and Future Work

This work studied dynamic pricing strategies in the streaming longitudinal data setting where the goal is to maximize the cumulative profit over time across a large number of customer segments. We proposed a pricing policy based on penalized stochastic gradient descent and provided regret bounds demonstrating the asymptotic optimality of the proposed policy. In particular, we showed that our PSGD algorithm controls the regret at the order of $\mathcal{O}(\sqrt{T})$




Figure 10: Regret plots of the proposed policy in Set-up 9 as b which governs the shift in the model parameters over time changes. The plots for $b = 0.5, 1, \infty$ are in red, green and blue respectively. In panels (a) Regret in original scale (b) $\log(\text{Regret})$ vs $\log(T)$. The dotted lines in panel (b) are the best fitted line with slope 0.5.

and that for any pricing policy, the Bayes regret cannot be of the lower order of $\mathcal{O}(\sqrt{T})$. Hence, as $T \rightarrow \infty$, the proposed algorithm is asymptotically rate-optimal. Our results show that for policy planning it is essential to incorporate available structural information as policies given by unshrunk models are highly sub-optimal.

There are several important future directions of our work. In future work, it will be useful to derive the regret of the proposed algorithm when the noise is non Gaussian, e.g. heavy-tailed distributions as such noise characteristics are often associated with observed demand data. Theoretically, it will be interesting to calculate the benefits of a batched version of the proposed algorithm 1 that is equipped for price exploration within segments though it might not be practically feasible due to spill-over effects. Also, here we have considered global spatial structure in the form of the SAR model in (7). In the future, it will be interesting to study the performance of PSGD in the presence of local shrinkage structures such as geographically weighted regression models (Fotheringham et al., 2003). Finally, it will be useful to evaluate the optimal regret in time-varying networks where the contiguity matrix $\mathbf{W}$ also changes over time.

Acknowledgments

This work is supported in part by an Adobe Faculty Research Award. Adel Javanmard is partially supported by the Sloan fellowship in mathematics, the NSF CAREER Award DMS-1844481, the NSF Award DMS-2311024 and an Amazon Faculty Research Award.

Appendix A. Organization of the Appendix

Here, we first present the proofs of results discussed in Section 3 of the main paper. The detailed proofs of the main results, Theorems 2 and 5, are provided in Section B of the appendix. All the other proofs of the intermediary results is provided in Section C of the Appendix.

Appendix B. Detailed Proofs of Theorem 2 and 5

B.1 Proof of Theorem 2

The total regret $\mathcal{B}(\boldsymbol{\theta}, \mathbf{p})$ can be written down as the sum of regrets over all the segments and time period:

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) = \sum_{t=1}^T \sum_{l=1}^L \mathcal{R}_{lt}.$$

Using Proposition 7 and Lemma 8, we can bound the regret as

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) = C_9 C_{10} \sum_{t=1}^T \sum_{l=1}^L n_{lt} \langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + C_9 C_{10} \sum_{t=1}^T \sum_{l=1}^L n_{lt} (p_{lt}(b_{lt} - \hat{b}_{lt}))^2.$$

Taking a maximum on the constants, we can rather bound the sum of the two terms $\sum_{t=1}^T \sum_{l=1}^L n_{lt} \langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2$ and $\sum_{t=1}^T \sum_{l=1}^L n_{lt} (p_{lt}(b_{lt} - \hat{b}_{lt}))^2$ to get a final bound on the regret.

We do the analysis for a fixed segment l first and then combine the regret across all the segments.

Lemma 9 *Consider model (12), true parameters $\mathbf{m}_{lt}$, b_{lt} and the output $\hat{\mathbf{m}}_{lt}$, $\hat{b}_{lt}$ from our PSGD pricing policy, the following holds with probability at least $1 - T^{-2}$*

$$\begin{aligned} & \sum_{t=1}^T n_{lt} \left(\langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + (p_{lt}(b_{lt} - \hat{b}_{lt}))^2 \right) \\ & \leq C_1 \sum_{t=1}^T \frac{1}{\eta_t} \|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 + C_2 \sum_{t=1}^T \frac{1}{\eta_t} |b_{l,t+1} - b_{lt}| + C_3 \sum_{t=1}^T \eta_t n_{lt}^2 + \frac{C_4}{\eta_{T+1}} + \mathcal{O}(\log T). \end{aligned}$$

With this lemma we have with probability at least $1 - L/T^2$,

$$\begin{aligned} & \sum_{t=1}^T \sum_{l=1}^L n_{lt} \left(\langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + (p_{lt}(b_{lt} - \hat{b}_{lt}))^2 \right) \\ & \leq C_1 \sum_{t=1}^T \sum_{l=1}^L \frac{1}{\eta_t} \|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 + C_2 \sum_{t=1}^T \sum_{l=1}^L \frac{1}{\eta_t} |b_{l,t+1} - b_{lt}| + C_3 \sum_{t=1}^T \sum_{l=1}^L \eta_t n_{lt}^2 + \frac{C_4 L}{\eta_{T+1}} + \mathcal{O}(\log T). \end{aligned}$$

Define the RHS as I .

Consider $\mathcal{G}$ to be the probabilistic event that the above is true, then $\mathbb{P}(\mathcal{G}^C) = L/T^2$.

Also, since the maximum price is M and we set a positive price, hence the maximum revenue lost on the event $\mathcal{G}^C$ is $\sum_{t=1}^T n_t M$. Assuming that $n_T^{\max} = \max_{t \leq T} n_t$, we have the maximum regret in the event $\mathcal{G}^C$ is $TMn_T^{\max}$.

The total regret is thus,

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) = \mathcal{B}(\boldsymbol{\theta}, \mathbf{p}|\mathcal{G}) + \mathcal{B}(\boldsymbol{\theta}, \mathbf{p}|\mathcal{G}^C) \leq I\mathbb{P}(\mathcal{G}) + Tn_T^{\max}\mathbb{P}(\mathcal{G}^C) \leq A + \frac{MLn_T^{\max}}{T}.$$

Since the last term is $\mathcal{O}(1/T)$, we have the required terms of the regret bound. $\blacksquare$

B.1.1 PROOF OF LEMMA 9

Let $\boldsymbol{\psi}_{lt} = (\mathbf{m}_{lt}, b_{lt})$ be the combined parameter space and $\mathbf{Q}_{lt} = (\mathbf{x}_{lt}, p_{lt})$ be the covariates and the price posted.

By Taylor expansion of the loss function we get, for some $\tilde{\boldsymbol{\psi}}_{lt}$ between $\hat{\boldsymbol{\psi}}_{lt}$ and $\boldsymbol{\psi}_{lt}$,

$$\mathcal{L}_{lt}(\hat{\boldsymbol{\psi}}_{lt}) - \mathcal{L}_{lt}(\boldsymbol{\psi}_{lt}) = \langle \nabla \mathcal{L}_{lt}(\hat{\boldsymbol{\psi}}_{lt}), \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt} \rangle - \frac{1}{2} \langle \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt}, \nabla^2 \mathcal{L}_{lt}(\tilde{\boldsymbol{\psi}}_{lt})(\hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt}) \rangle. \quad (30)$$

Simplifying the loss function in terms of $\boldsymbol{\psi}_{lt}$ and $\mathbf{Q}_{lt}$ gives us

$$\mathcal{L}_{lt}(\boldsymbol{\psi}) = -(y_{lt} \log \Phi(\mathbf{Q}_{lt}\boldsymbol{\psi}) + \tilde{y}_{lt} \log \Phi(-\mathbf{Q}_{lt}\boldsymbol{\psi})),$$

where $\tilde{y}_{lt} = n_{lt} - y_{lt}$. The second derivative of the loss function can thus be computed as

$$\nabla^2 \mathcal{L}_{lt}(\boldsymbol{\psi}) = - \left(y_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=\mathbf{Q}_{lt}\boldsymbol{\psi}} + \tilde{y}_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=-\mathbf{Q}_{lt}\boldsymbol{\psi}} \right).$$

Let $c_{\mathcal{L}} = \min \left\{ -\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=\mathbf{Q}_{lt}\boldsymbol{\psi}}, -\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=-\mathbf{Q}_{lt}\boldsymbol{\psi}} \right\}$. Based on our assumptions, $\mathbf{Q}_{lt}$ and $\boldsymbol{\psi}$ are bounded and so there exists c such that $|\mathbf{Q}_{lt}\boldsymbol{\psi}| < c$. Since Φ is log-concave hence the second derivative is negative and $-\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta) > 0$. Particularly since the second derivative only approaches 0 when ζ goes to ∞ or $-\infty$, hence on the bounded set with $|\zeta| < c$, second derivative is bounded away from zero implying $c_{\mathcal{L}} > 0$.

Then,

$$\begin{aligned} \nabla^2 \mathcal{L}_{lt}(\boldsymbol{\psi}) &= - \left(y_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=\mathbf{Q}_{lt}\boldsymbol{\psi}} + \tilde{y}_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=-\mathbf{Q}_{lt}\boldsymbol{\psi}} \right) \\ &= \left(y_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \left(-\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=\mathbf{Q}_{lt}\boldsymbol{\psi}} \right) + \tilde{y}_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \left(-\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=-\mathbf{Q}_{lt}\boldsymbol{\psi}} \right) \right) \\ &\geq (y_{lt} + \tilde{y}_{lt}) \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T \min \left\{ -\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=\mathbf{Q}_{lt}\boldsymbol{\psi}}, -\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta)|_{\zeta=-\mathbf{Q}_{lt}\boldsymbol{\psi}} \right\} \\ &= n_{lt} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T c_{\mathcal{L}}. \end{aligned}$$

Where the last equality follows since $y_{lt} + \tilde{y}_{lt} = n_{lt}$.

Using this in (30)

$$\begin{aligned}
 \mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_t) &\leq \langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{lt} - \psi_t \rangle - \frac{1}{2} \langle \hat{\psi}_{lt} - \psi_t, n_{lt} c_{\mathcal{L}} \mathbf{Q}_{lt} \mathbf{Q}_{lt}^T (\hat{\psi}_{lt} - \psi_t) \rangle \\
 &= \langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{lt} - \psi_t \rangle - \frac{n_{lt} c_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_t, \mathbf{Q}_{lt} \rangle^2 \\
 &= \langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{l,t+1} - \psi_t \rangle + \langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{lt} - \hat{\psi}_{l,t+1} \rangle - \frac{n_{lt} c_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_t, \mathbf{Q}_{lt} \rangle^2.
 \end{aligned} \tag{31}$$

Our update rules in (19), gives us

$$\begin{aligned}
 \hat{b}_{l,t+1} &= \Pi_{\Theta_b}(\hat{b}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^b), \\
 \hat{\mathbf{m}}_{l,t+1} &= \Pi_{\Theta_m}(\hat{\mathbf{m}}_{lt} - \eta_t \nabla \mathcal{L}_{lt}^m).
 \end{aligned}$$

The updates defined are common OMD updates and can be rewritten as

$$\hat{\psi}_{l,t+1} = \arg \min_{\psi} \langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \psi \rangle + \frac{1}{2\eta_t} \|\psi - \hat{\psi}_{lt}\|^2.$$

Since the above loss function is convex and $\hat{\psi}_{l,t+1}$ is the minimizer, we get

$$\langle \psi - \hat{\psi}_{l,t+1}, \eta_t \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}) + \hat{\psi}_{l,t+1} - \hat{\psi}_{lt} \rangle \geq 0.$$

Putting $\psi = \psi_t$ above, we get $\langle \hat{\psi}_{l,t+1} - \psi_t, \eta_t \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}) \rangle \leq \langle \psi_t - \hat{\psi}_{l,t+1}, \hat{\psi}_{l,t+1} - \hat{\psi}_{lt} \rangle$. Also, note that

$$\langle \psi_t - \hat{\psi}_{l,t+1}, \hat{\psi}_{l,t+1} - \hat{\psi}_{lt} \rangle = \frac{1}{2} \left(\|\psi_t - \hat{\psi}_{lt}\|^2 - \|\psi_t - \hat{\psi}_{l,t+1}\|^2 - \|\hat{\psi}_{l,t+1} - \hat{\psi}_{lt}\|^2 \right).$$

With the above two equations the first term in (31) is bounded as:

$$\langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{l,t+1} - \psi_t \rangle \leq \frac{1}{2\eta_t} \left(\|\psi_t - \hat{\psi}_{lt}\|^2 - \|\psi_t - \hat{\psi}_{l,t+1}\|^2 - \|\hat{\psi}_{l,t+1} - \hat{\psi}_{lt}\|^2 \right).$$

Using the inequality $ab \leq (a^2 + b^2)/2$, the second term in (31) can be bounded as,

$$\langle \nabla \mathcal{L}_{lt}(\hat{\psi}_{lt}), \hat{\psi}_{lt} - \hat{\psi}_{l,t+1} \rangle \leq \frac{1}{2\eta_t} \|\hat{\psi}_{lt} - \hat{\psi}_{l,t+1}\|^2 + \frac{\eta_t}{2} \|\nabla \mathcal{L}_{lt}(\hat{\psi}_{lt})\|^2. \tag{32}$$

Also, $\nabla \mathcal{L}_{lt}(\psi) = - \left(y_{lt} \mathbf{Q}_{lt} \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta) |_{\zeta=\mathbf{Q}_{lt}\psi} - \tilde{y}_{lt} \mathbf{Q}_{lt} \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta) |_{\zeta=-\mathbf{Q}_{lt}\psi} \right)$.

Let $C_{\mathcal{L}} = \max\{-\frac{\partial}{\partial \zeta^2} \log \Phi(\zeta) |_{\zeta=\mathbf{Q}_{lt}\psi}, \frac{\partial}{\partial \zeta^2} \log \Phi(\zeta) |_{\zeta=-\mathbf{Q}_{lt}\psi}\}$ in the restricted space. Hence, $\|\nabla \mathcal{L}_{lt}(\hat{\psi}_{lt})\|^2 \leq C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2$.

Combining all the parts we have,

$$\mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_t) \leq \frac{1}{2\eta_t} \|\psi_t - \hat{\psi}_{lt}\|^2 - \frac{1}{2\eta_t} \|\psi_t - \hat{\psi}_{l,t+1}\|^2 + \frac{\eta_t}{2} C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2 - \frac{n_{lt} c_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_t, \mathbf{Q}_{lt} \rangle^2.$$

Adding and subtracting $\|\psi_{l,t+1} - \hat{\psi}_{l,t+1}\|^2$ to above we get

$$\begin{aligned} \mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_t) &\leq \frac{1}{2\eta_t} \left(\|\psi_{lt} - \hat{\psi}_{lt}\|^2 - \|\psi_{l,t+1} - \hat{\psi}_{l,t+1}\|^2 \right) \\ &\quad + \frac{1}{2\eta_t} \left(\|\psi_{l,t+1} - \hat{\psi}_{l,t+1}\|^2 - \|\psi_{lt} - \hat{\psi}_{l,t+1}\|^2 \right) \\ &\quad + \frac{\eta_t}{2} C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2 - \frac{n_{lt} C_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_{lt}, \mathbf{Q}_{lt} \rangle^2. \end{aligned} \quad (33)$$

The second term can be simplified as

$$\|\psi_{l,t+1} - \hat{\psi}_{l,t+1}\|^2 - \|\psi_{lt} - \hat{\psi}_{l,t+1}\|^2 = \langle \psi_{l,t+1} + \psi_{lt} - 2\hat{\psi}_{l,t+1}, \psi_{l,t+1} - \psi_{lt} \rangle \leq 4C_{\psi} \|\psi_{l,t+1} - \psi_{lt}\|_2,$$

where C_{ψ} is $\max \|\psi\|$ and $C_{\psi} \leq 2C_b + 2C_m$, since $\psi = (\mathbf{m}, b)$.

Summing both sides of (33) over $t = 1, \dots, T$, we get $\sum_{t=1}^T (\mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_t))$ is bounded above by:

$$\begin{aligned} &\frac{\|\psi_{l1} - \hat{\psi}_{l1}\|^2}{2\eta_1} + \sum_{t=2}^T \|\psi_{lt} - \hat{\psi}_{lt}\|^2 \left(\frac{1}{2\eta_{t+1}} - \frac{1}{2\eta_t} \right) + 4C_{\psi} \sum_{t=1}^T \frac{1}{2\eta_t} \|\psi_{l,t+1} - \psi_{lt}\|_2 \\ &+ \sum_{t=1}^T \frac{\eta_t}{2} C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2 - \sum_{t=1}^T \frac{n_{lt} C_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_{lt}, \mathbf{Q}_{lt} \rangle^2. \end{aligned}$$

Under the assumption that η_t are non-decreasing,

$$\frac{\|\psi_{l1} - \hat{\psi}_{l1}\|^2}{2\eta_1} + \sum_{t=2}^T \|\psi_{lt} - \hat{\psi}_{lt}\|^2 \left(\frac{1}{2\eta_{t+1}} - \frac{1}{2\eta_t} \right) \leq \frac{4C_{\psi}^2}{2\eta_1} + 4C_{\psi}^2 \sum_{t=2}^T \left(\frac{1}{2\eta_{t+1}} - \frac{1}{2\eta_t} \right) = \frac{4C_{\psi}^2}{2\eta_{T+1}}.$$

Hence, we finally have

$$\begin{aligned} \sum_{t=1}^T (\mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_t)) &\leq \frac{4C_{\psi}^2}{2\eta_{T+1}} + 4C_{\psi} \sum_{t=1}^T \frac{1}{2\eta_t} \|\psi_{l,t+1} - \psi_{lt}\|_2 \\ &\quad + \sum_{t=1}^T \frac{\eta_t}{2} C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2 - \sum_{t=1}^T \frac{n_{lt} C_{\mathcal{L}}}{2} \langle \hat{\psi}_{lt} - \psi_{lt}, \mathbf{Q}_{lt} \rangle^2. \end{aligned}$$

Define

$$A := \frac{4C_{\psi}^2}{2\eta_{T+1}} + 4C_{\psi} \sum_{t=1}^T \frac{1}{2\eta_t} \|\psi_{l,t+1} - \psi_{lt}\|_2 + \sum_{t=1}^T \frac{\eta_t}{2} C_{\mathcal{L}}^2 n_{lt}^2 \|\mathbf{Q}_{lt}\|^2. \quad (34)$$

Since, $\|\psi_{l,t+1} - \psi_{lt}\|_2 \leq 2(\|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 + |b_{l,t+1} - b_{lt}|)$ and $\|\mathbf{Q}_{lt}\|^2$ is bounded, we can simplify A as

$$A := \tilde{C}_1 \sum_{t=1}^T \frac{1}{\eta_t} \|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 + \tilde{C}_2 \sum_{t=1}^T \frac{1}{\eta_t} |b_{l,t+1} - b_{lt}| + \tilde{C}_3 \sum_{t=1}^T \eta_t n_{lt}^2 + \frac{\tilde{C}_4}{\eta_{T+1}}.$$

Note that in order to prove the lemma we need to show a bound on $\sum_{t=1}^T n_{lt} \left(\langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + (p_{lt}(b_{lt} - \hat{b}_{lt}))^2 \right)$ which is same as showing a bound on $\sum_{t=1}^T n_{lt} \langle \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt}, \mathbf{Q}_{lt} \rangle^2$, since $\boldsymbol{\psi}_{lt} = (\mathbf{m}_{lt}, b_{lt})$ and $\mathbf{Q}_{lt} = (\mathbf{x}_{lt}, p_{lt})$.

We next provide a lower bound on the cumulative difference $\sum_{t=1}^T (\mathcal{L}_{lt}(\hat{\boldsymbol{\psi}}_{lt}) - \mathcal{L}_{lt}(\boldsymbol{\psi}_{lt}))$. Write

$$\mathcal{L}_{ltk}(\boldsymbol{\psi}_{lt}) - \mathcal{L}_{ltk}(\hat{\boldsymbol{\psi}}_{lt}) \leq \langle \nabla \mathcal{L}_{ltk}(\boldsymbol{\psi}_{lt}), \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt} \rangle := D_{tk}, \quad (35)$$

using convexity of the loss $\mathcal{L}_{ltk}$. We also have

$$\begin{aligned} \nabla \mathcal{L}_{ltk}(\boldsymbol{\psi}) &= -(y_{ltk} \frac{\partial}{\partial \boldsymbol{\psi}} \log \Phi(\mathbf{Q}_{lt} \boldsymbol{\psi}) - (1 - y_{ltk}) \frac{\partial}{\partial \boldsymbol{\psi}} \log \Phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})) \\ &= \mathbf{Q}_{lt} \left(-y_{ltk} \frac{\phi(\mathbf{Q}_{lt} \boldsymbol{\psi})}{\Phi(\mathbf{Q}_{lt} \boldsymbol{\psi})} + (1 - y_{ltk}) \frac{\phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})}{\Phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})} \right). \end{aligned}$$

Let $\mathcal{F}_t$ be the σ -field generated by the noise till time t . Then, since $\hat{\boldsymbol{\psi}}_{lt}$ only depends on noise till time t , $\mathbb{E}[D_{tk} | \mathcal{F}_{t-1}] = \langle \mathbb{E}[\nabla \mathcal{L}_{ltk} | \mathcal{F}_{t-1}], \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt} \rangle$. In addition, $\mathbb{E}[\nabla \mathcal{L}_{ltk} | \mathcal{F}_{t-1}] = 0$ using the fact that $\mathbb{P}(Y_{ltk} = 1) = \Phi(\mathbf{Q}_{lt} \boldsymbol{\psi})$ and $\mathbb{P}(Y_{ltk} = 0) = \Phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})$. Therefore, the partial sums of D_{tk} is a martingale with respect to the filtration $\mathcal{F}_t$.

Also, as described above $\left(-y_{ltk} \frac{\phi(\mathbf{Q}_{lt} \boldsymbol{\psi})}{\Phi(\mathbf{Q}_{lt} \boldsymbol{\psi})} + (1 - y_{ltk}) \frac{\phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})}{\Phi(-\mathbf{Q}_{lt} \boldsymbol{\psi})} \right)$ is bounded above with $C_{\mathcal{L}}$. Hence $|D_{tk}| \leq \beta_t := C_{\mathcal{L}} |\langle \mathbf{Q}_{lt}, \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt} \rangle|$. Using convexity of $e^{\lambda z}$, for any $\lambda \in \mathbb{R}$ we have

$$\begin{aligned} \mathbb{E} \left[e^{\lambda D_{tk}} \mid \mathcal{F}_{t-1} \right] &\leq \mathbb{E} \left[\frac{\beta_t - D_{tk}}{2\beta_t} e^{-\lambda \beta_t} + \frac{\beta_t + D_{tk}}{2\beta_t} e^{\lambda \beta_t} \mid \mathcal{F}_{t-1} \right] \\ &= \mathbb{E} \left[\frac{e^{-\lambda \beta_t} + e^{\lambda \beta_t}}{2} \right] + \mathbb{E}[D_{tk} \mid \mathcal{F}_{t-1}] \left(\frac{e^{-\lambda \beta_t} + e^{\lambda \beta_t}}{2\beta_t} \right) = \cosh(\lambda \beta_t) \leq e^{\lambda^2 \beta_t^2 / 2}. \end{aligned}$$

where $\beta_t = C_{\mathcal{L}} |\langle \mathbf{Q}_{lt}, \hat{\boldsymbol{\psi}}_{lt} - \boldsymbol{\psi}_{lt} \rangle|$. We next use the following result from (Javanmard, 2017, Proposition C.1).

Proposition 10 (Javanmard, 2017, Proposition C.1) *Consider a martingale difference sequence D_t adapted to a filtration $\mathcal{F}_t$, such that for any $\lambda \geq 0$, $\mathbb{E}[e^{\lambda D_t} \mid \mathcal{F}_{t-1}] \leq e^{\lambda^2 \sigma_t^2 / 2}$. Then, for $D(T) = \sum_{t=1}^T D_t$, the following holds true:*

$$\mathbb{P}(D(T) \geq \xi) \leq e^{-\xi^2 / (2 \sum_{t=1}^T \sigma_t^2)}.$$

We apply the above theorem with $D(T) = \sum_{t=1}^T \sum_{k=1}^{n_{lt}} D_{tk}$. Invoking (35), this gives us,

$$\mathbb{P} \left(\sum_{t=1}^T (\mathcal{L}_{lt}(\hat{\boldsymbol{\psi}}_{lt}) - \mathcal{L}_{lt}(\boldsymbol{\psi}_{lt})) \leq -2C_{\mathcal{L}} \sqrt{\log T} \left\{ \sum_{t=1}^T n_{lt} \langle \mathbf{Q}_{lt}, \boldsymbol{\psi}_{lt} - \hat{\boldsymbol{\psi}}_{lt} \rangle^2 \right\}^{1/2} \right) \leq \frac{1}{T^2}.$$

Hence with probability at least $1 - 1/T^2$,

$$\sum_{t=1}^T (\mathcal{L}_{lt}(\hat{\boldsymbol{\psi}}_{lt}) - \mathcal{L}_{lt}(\boldsymbol{\psi}_{lt})) \geq -2C_{\mathcal{L}} \sqrt{\log T} \left\{ \sum_{t=1}^T n_{lt} \langle \mathbf{Q}_{lt}, \boldsymbol{\psi}_{lt} - \hat{\boldsymbol{\psi}}_{lt} \rangle^2 \right\}^{1/2}.$$

Let $B = \sum_{t=1}^T n_{lt} \langle \hat{\psi}_{lt} - \psi_{lt}, \mathbf{Q}_{lt} \rangle^2$, then with the complete analysis till now we have

$$-2C_{\mathcal{L}}\sqrt{B \log T} \leq \sum_{t=1}^T \left(\mathcal{L}_{lt}(\hat{\psi}_{lt}) - \mathcal{L}_{lt}(\psi_{lt}) \right) \leq A - \frac{c_{\mathcal{L}}}{2} B.$$

Hence, $B - (4C_{\mathcal{L}}/c_{\mathcal{L}})\sqrt{B \log T} \leq (2/c_{\mathcal{L}})A$. Consider two cases:

Case 1: $\sqrt{B \log T} \leq (c_{\mathcal{L}}/8C_{\mathcal{L}})B$, then $B \leq 4A/c_{\mathcal{L}}$.

Case 2: $\sqrt{B \log T} \geq (c_{\mathcal{L}}/8C_{\mathcal{L}})B$, then $B \leq (8C_{\mathcal{L}}/c_{\mathcal{L}})^2 \log T$.

Combining the two cases, we have

$$B \leq \frac{4A}{c_{\mathcal{L}}} + \mathcal{O}(\log T).$$

Substituting for

$$B = \sum_{t=1}^T n_{lt} \langle \hat{\psi}_{lt} - \psi_{lt}, \mathbf{Q}_{lt} \rangle^2 = \sum_{t=1}^T n_{lt} \left(\langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + p_{lt}^2 (b_{lt} - \hat{b}_{lt})^2 \right),$$

and A from (34) we obtain the desired result. ■

B.2 Proof of Theorem 5

Recall the variance of segment l in the utility model given by $V_{lt}^2 = \|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 \tau^2 + \sigma^2$. We assume that we have a known fixed ρ , the auto-correlation parameter, and so the variances do not change over time.

Indicate the variances of segments by $V_1, V_2, \dots, V_L$. Our utility model is thus

$$\tilde{U}_{ltk} = \frac{\beta_t}{V_l} p_{lt} + \mathbf{x}'_{lt} \frac{\boldsymbol{\mu}_t}{V_l} + Z_{ltk}.$$

Without loss of generality, assume that $\mathbf{x}_{lt}$ is of dimension one. We would give a small variation that would work for any dimension as well. Let $v_1, v_2, \dots, v_L$ be the inverse of the fixed variances. The model is thus,

$$\tilde{U}_{ltk} = \beta_t v_l p_{lt} + x_{lt} \mu_t v_l + Z_{ltk}.$$

Assume that $-\beta_t = \mu_t = \gamma$, i.e. the parameters do not change over time and are negative of each other. In this setup $U_{ltk}^0 = v_l \gamma (x_{lt} - p_{lt})$, the noiseless utility. Here setting $p_{lt} = x_{lt}$ would be uninformative since we would just observe noise, and we cannot get any information about the unknown parameter γ . In addition, in our model a price p_{lt}^* is optimum if it satisfies

$$p_{lt}^* = -\frac{1}{\beta_t v_l} \frac{\Phi(\beta_t v_l p_{lt}^* + x_{lt} \mu_t v_l)}{\phi(\beta_t v_l p_{lt}^* + x_{lt} \mu_t v_l)}.$$

Under the assumption that $-\beta_t = \mu_t = \gamma$, this reduces to

$$p_{lt}^* = \frac{1}{\gamma v_l} \frac{\Phi(v_l \gamma (x_{lt} - p_{lt}^*))}{\phi(v_l \gamma (x_{lt} - p_{lt}^*))}.$$

Therefore, for $\gamma_0 := (v_l x_{lt})^{-1} \Phi(0)/\phi(0)$, the uninformative price is optimal prices, i.e., $p_{lt}^*(\gamma_0) = x_{lt}$.

Note that if x_{lt} was of higher dimension, we could set $\mathbf{x}_{lt} = (a/d, a/d, a/d, \dots, a/d)$ with d the dimension of $\mathbf{x}_{lt}$ and set $\boldsymbol{\mu}_t = (\gamma, \gamma, \gamma, \dots, \gamma)$ to get the exactly same result: $p_{lt}^*(\gamma_0) = a$ for $\gamma_0 = (v_l a)^{-1} \Phi(0)/\phi(0)$ is an uninformative price.

Now that we know the existence of a setting where the uninformative prices are optimal prices, we can show that the regret is at least of the order of $\sqrt{T}$.

We construct a problem class $(\Gamma, \{\mathcal{P}_{lt}\})$, for $l = 1, \dots, L$, $t = 1, \dots, T$ as follows. Recall γ_0 the parameter for which the optimal price is uninformative. We use the shorthand $r_{lt}(p, \gamma)$ to denote the expected revenue obtained from a typical customer from segment l at time t , if the model parameter is γ . Therefore, recalling our utility model $U_{ltk} = v_l \gamma(x_{lt} - p_{lt}) + Z_{ltk}$, we have $r_{lt}(p, \gamma) = p(\Phi(v_l \gamma(x_{lt} - p)))$. By optimality of $p_{lt}^*(\gamma_0)$, we have $r_{lt}''(p_{lt}^*(\gamma_0), \gamma_0) < -2c$ for some constant $c > 0$, and by continuity of r_{lt}'' we can find a neighborhood $\mathcal{P}_{lt}$ around $p_{lt}^*(\gamma_0)$ such that $r_{lt}''(p, \gamma_0) < -c$ for all $p \in \mathcal{P}_{lt}$. We next consider the mapping $\gamma \mapsto p_{lt}^*(\gamma)$. By continuity of this mapping, we can find a small enough neighborhood Γ_{lt} around γ_0 such that the optimal prices $p_{lt}^*(\gamma) \in \mathcal{P}_{lt}$ for all $\gamma \in \Gamma_{lt}$. Finally, we take $\Gamma := \cap_{t=1}^T \cap_{l=1}^L \Gamma_{lt}$. Note that Γ is non-empty because $\gamma_0 \in \Gamma$. Furthermore, by our construction we have the following properties for the problem class $(\Gamma, \{\mathcal{P}\}_{lt})$, for $l = 1, \dots, L$ and $t = 1, \dots, T$:

- For all $\gamma \in \Gamma$, we have $p_{lt}^*(\gamma) \in \mathcal{P}_{lt}$.
- For all prices $p \in \mathcal{P}_{lt}$, we have $r_{lt}''(p, \gamma_0) < -c$.

For any pricing policy π and a parameter $\gamma \in \Gamma$, let $f_t^{\pi, \gamma} : \{0, 1\}^{N_t} \rightarrow [0, 1]$ be the probability distribution function for all the consumers purchase responses $\mathbf{Y} = (Y_{ljk}, \ell = 1, \dots, L, j = 1, \dots, t, k = 1, \dots, n_{lj})$ until time t . Here, $N_t = \sum_{j=1}^t \sum_{l=1}^L n_{lj}$, under policy π and model parameter γ . The pricing policy uses all the sales data till time $t-1$ to give a price p_{lt}^* . We use $\mathbf{y}_t \in \{0, 1\}^{N_t}$ to denote all sales data till time t . So, if the pricing policy gives the prices $p_{lt} := \pi(\mathbf{y}_{t-1})$ for all the time periods, then

$$f_t^{\pi, \gamma}(\mathbf{y}_t) = \prod_{j=1}^t \prod_{l=1}^L \prod_{k=1}^{n_{lj}} q_{lj}(p_{lj}, \gamma)^{y_{lj k}} (1 - q_{lj}(p_{lj}, \gamma))^{1 - y_{lj k}},$$

where $q_{lj}(p_{lj}, \gamma) = \Phi(v_l \gamma(x_{lj} - p_{lj}))$. We next want to show that for γ_0 , the parameter for which the uninformative price is optimal, any policy incurs a large regret if it tries to learn γ_0 . Formally, we aim to show that

$$\mathcal{R}_t^\pi(\gamma_0) \geq C \frac{1}{(\gamma_0 - \gamma)^2} \text{KL}(f_t^{\pi, \gamma_0}, f_t^{\pi, \gamma}).$$

We employ the chain rule for KL divergence (Cover and Thomas, 1991),

$$\begin{aligned}
 \text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}) &= \sum_{s=1}^t \text{KL}(f_s^{\pi, \gamma_0}; f_s^{\pi, \gamma} | \mathbf{Y}_{s-1}) \\
 &= \sum_{s=1}^t \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_s^{\pi, \gamma_0}(\mathbf{y}_{s-1}) \log \left(\frac{f_s^{\pi, \gamma_0}(y_{lsk} | \mathbf{y}_{s-1})}{f_s^{\pi, \gamma}(y_{lsk} | \mathbf{y}_{s-1})} \right) \\
 &= \sum_{s=1}^t \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_{s-1}^{\pi, \gamma_0}(\mathbf{y}_{s-1}) \sum_{l=1}^L \sum_{k=1}^{n_{ls}} \sum_{y_{lsk} \in \{0,1\}} f_s^{\pi, \gamma_0}(y_{lsk} | \mathbf{y}_{s-1}) \log \left(\frac{f_s^{\pi, \gamma_0}(y_{lsk} | \mathbf{y}_{s-1})}{f_s^{\pi, \gamma}(y_{lsk} | \mathbf{y}_{s-1})} \right) \\
 &= \sum_{s=1}^t \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_{s-1}^{\pi, \gamma_0}(\mathbf{y}_{s-1}) \sum_{l=1}^L \sum_{k=1}^{n_{ls}} \text{KL} \left(f_s^{\pi, \gamma_0}(y_{lsk} | \mathbf{y}_{s-1}); f_s^{\pi, \gamma}(y_{lsk} | \mathbf{y}_{s-1}) \right).
 \end{aligned}$$

Based on the definition of $f_t^{\pi, \gamma}$, $f_s^{\pi, \gamma_0}(y_{lsk})$ is distributed as Bernoulli $q_{ls}(p_{ls}, \gamma_0)$ and $f_s^{\pi, \gamma}(y_{lsk})$ is distributed as Bernoulli $q_{ls}(p_{ls}, \gamma)$. Using the fact that for Bernoulli random variables $B_1 \sim \text{Bern}(q_1)$, $B_2 \sim \text{Bern}(q_2)$, we have $\text{KL}(B_1, B_2) \leq \frac{(q_1 - q_2)^2}{q_2(1 - q_2)}$, we get

$$\begin{aligned}
 \text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}) &\leq \sum_{s=1}^t \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_{s-1}^{\pi, \gamma_0}(\mathbf{y}_{s-1}) \sum_{l=1}^L \sum_{k=1}^{n_{ls}} \frac{(q_{ls}(p_{ls}, \gamma_0) - q_{ls}(p_{ls}, \gamma))^2}{q_{ls}(p_{ls}, \gamma)(1 - q_{ls}(p_{ls}, \gamma))} \\
 &= \sum_{s=1}^t \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_{s-1}^{\pi, \gamma_0}(\mathbf{y}_{s-1}) \sum_{l=1}^L n_{ls} \frac{(q_{ls}(p_{ls}, \gamma_0) - q_{ls}(p_{ls}, \gamma))^2}{q_{ls}(p_{ls}, \gamma)(1 - q_{ls}(p_{ls}, \gamma))}. \quad (36)
 \end{aligned}$$

Since the prices and the parameters are bounded, and q_{lt} is the normal distribution function, q_{lt} is bounded away from zero. Hence, there exists constant C such that $q_{ls}(1 - q_{ls}) \geq C$.

Also, $q_{ls}(p_{ls}^*, \gamma) = \Phi(v_l \gamma(x_{ls} - p_{ls}^*))$ and $q_{ls}(p_{ls}^*, \gamma_0) = \Phi(v_l \gamma_0(x_{ls} - p_{ls}^*))$. Since we are working on a bounded set, the distribution function Φ is Lipschitz as well. Hence,

$$\begin{aligned}
 q_{ls}(p_{ls}, \gamma_0) - q_{ls}(p_{ls}, \gamma) &= \Phi(v_l \gamma_0(x_{ls} - p_{ls})) - \Phi(v_l \gamma(x_{ls} - p_{ls})) \\
 &\leq C(v_l \gamma_0(x_{ls} - p_{ls}) - v_l \gamma(x_{ls} - p_{ls})) \\
 &= C v_l (\gamma_0 - \gamma)(x_{ls} - p_{ls}) \\
 &= C v_l (\gamma_0 - \gamma)(p_{ls}^*(\gamma_0) - p_{ls}),
 \end{aligned}$$

where $p_{ls}^*(\gamma_0)$ is the optimal price for when the parameter is γ_0 . Recall that by the definition of γ_0 , the optimal price for γ_0 is x_{lt} . We thus have

$$(q_{ls}(p_{ls}, \gamma_0) - q_{ls}(p_{ls}, \gamma))^2 \leq C(\gamma_0 - \gamma)^2 (p_{ls}^*(\gamma_0) - p_{ls})^2.$$

Using the above bound in (36), we get

$$\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}) \leq C(\gamma - \gamma_0)^2 \sum_{s=1}^t \sum_{l=1}^L n_{ls} \sum_{\mathbf{y}_{s-1} \in \{0,1\}^{N_{s-1}}} f_{s-1}^{\pi, \gamma_0}(\mathbf{y}_{s-1}) (p_{ls}^*(\gamma_0) - p_{ls})^2.$$

The inner summation is indeed the expectation with respect to γ_0 , by noting that p_{ls} is a measurable function of $\mathbf{y}_{s-1}$. Hence, we have

$$\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}) \leq C(\gamma - \gamma_0)^2 \sum_{s=1}^t \sum_{l=1}^L n_{ls} \mathbb{E}_{\gamma_0} (p_{ls}^*(\gamma_0) - p_{ls})^2. \quad (37)$$

By the construction of problem class $(\Gamma, \{\mathcal{P}_{lt}\})$, we have $r_{lt}''(p, \gamma_0) \leq -c$, for $\gamma \in \Gamma$ and $p \in \mathcal{P}_{lt}$. Therefore, by Taylor expansion of $r_{ls}(p, \gamma)$ around p_{ls}^* , we obtain

$$r_{ls}(p_{ls}, \gamma_0) = r_{ls}(p_{ls}^*(\gamma_0), \gamma_0) + r_{ls}'(p_{ls}^*(\gamma_0), \gamma_0)(p_{ls} - p_{ls}^*(\gamma_0)) + \frac{1}{2} r_{ls}''(\tilde{p}, \gamma_0)(p_{ls} - p_{ls}^*(\gamma_0))^2,$$

for some $\tilde{p}$ between p_{ls} and p_{ls}^* . By optimality of p_{ls}^* we have $r_{ls}'(p_{ls}^*(\gamma_0), \gamma_0) = 0$. In addition, since $\tilde{p} \in \mathcal{P}_{ls}$, we have $r_{ls}''(\tilde{p}, \gamma_0) < -c$, which implies that

$$(p_{ls} - p_{ls}^*(\gamma_0))^2 \leq \frac{2}{c} \left(r_{ls}(p_{ls}^*(\gamma_0), \gamma_0) - r_{ls}(p_{ls}, \gamma_0) \right).$$

Using the above bound in (37), we arrive at

$$\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}) \leq C(\gamma - \gamma_0)^2 \sum_{s=1}^t \sum_{l=1}^L n_{ls} \mathbb{E}_{\gamma_0} [r_{ls}(p_{ls}^*(\gamma_0), \gamma_0) - r_{ls}(p_{ls}, \gamma_0)] \leq C(\gamma - \gamma_0)^2 \text{Reg}_t, \quad (38)$$

which completes the proof (28).

We next proceed with our proof for bound (29). Recall the optimality condition

$$v_l \gamma p_{lt}^*(\gamma) = \frac{\Phi(v_l \gamma (x_{lt} - p_{lt}^*(\gamma)))}{\phi(v_l \gamma (x_{lt} - p_{lt}^*(\gamma)))}.$$

Differentiating with respect to γ on both sides we get,

$$v_l p_{lt}^*(\gamma) + v_l \gamma \frac{d}{d\gamma} p_{lt}^*(\gamma) = v_l \left(x_{lt} - p_{lt}^*(\gamma) - \gamma \frac{d}{d\gamma} p_{lt}^*(\gamma) \right) \kappa(\gamma),$$

where

$$\kappa(\gamma) = \frac{\phi^2(v_l \gamma (x_{lt} - p_{lt}^*(\gamma))) - \Phi(v_l \gamma (x_{lt} - p_{lt}^*(\gamma))) \phi'(v_l \gamma (x_{lt} - p_{lt}^*(\gamma)))}{\phi^2(v_l \gamma (x_{lt} - p_{lt}^*(\gamma)))}.$$

By rearranging the terms we have

$$\frac{d}{d\gamma} p_{lt}^*(\gamma) = \frac{1}{\gamma} \left(-p_{lt}^*(\gamma) + \frac{k(\gamma)}{1 + k(\gamma)} \right).$$

Since we are working on finite sets, we can restrict the problem class Γ , such that $|\frac{d}{d\gamma} p(\gamma)| > C$, for some constant C and all $\gamma \in \Gamma$. Therefore, by an application of the Mean Value Theorem, we have

$$|p_{lt}^*(\gamma) - p_{lt}^*(\gamma_0)| \geq C|\gamma - \gamma_0|.$$

Let $\gamma_1 := \gamma_0 + 1/(4T^{1/4})$. Using the above bound, the optimal prices for γ_0 and γ_1 are apart by at least $C/(4T^{1/4})$.

Consider two disjoint sets D_1 and D_0 of prices, as follows:

$$D_{\gamma_0} := \left\{ p : |p - p_{lt}^*(\gamma_0)| \leq \frac{C}{10T^{1/4}} \right\}, \quad D_{\gamma_1} := \left\{ p : |p - p_{lt}^*(\gamma_1)| \leq \frac{C}{10T^{1/4}} \right\}.$$

Note that D_{γ_0} and D_{γ_1} are disjoint since $|p_{lt}^*(\gamma_1) - p_{lt}^*(\gamma_0)| \geq C/(4T^{1/4})$.

For $\gamma \in \{\gamma_0, \gamma_1\}$, if the posted price p_{lt} is not in the set D_γ , then the instantaneous regret is at least

$$r_{lt}(p_{lt}^*(\gamma), \gamma) - r_{lt}(p_{lt}, \gamma) \geq \frac{c}{2}(p_{lt}^*(\gamma) - p_{lt}) \geq \left(\frac{cC}{20}\right)^2 \frac{1}{\sqrt{T}}.$$

Hence following a similar proof strategy as in (Broder and Rusmevichientong, 2012, Lemma 3.4), we have

$$\text{Reg}_T^{\pi, \gamma_0} + \text{Reg}_T^{\pi, \gamma_1} \geq \left(\frac{cC}{20}\right)^2 \frac{1}{\sqrt{T}} \sum_{t=1}^T \sum_{l=1}^L n_{lt} (\mathbb{P}_{\gamma_0}(p_{lt} \notin D_{\gamma_0}) + \mathbb{P}_{\gamma_1}(p_{lt} \notin D_{\gamma_1})).$$

Note that p_{lt} is measurable with respect to measure $f_{t-1}^{\pi, \gamma}$ under the model γ . Therefore, by using a standard result on the minimum error in a simple hypothesis test (Tsybakov, 2004, Theorem 2.2), we have

$$\begin{aligned} \text{Reg}_T^{\pi, \gamma_0} + \text{Reg}_T^{\pi, \gamma_1} &\geq \frac{C_1}{\sqrt{T}} \sum_{t=1}^T \sum_{l=1}^L n_{lt} e^{-\text{KL}(f_{t-1}^{\pi, \gamma_0}; f_{t-1}^{\pi, \gamma_1})} \\ &\geq \frac{C_1}{\sqrt{T}} N_T e^{-\text{KL}(f_T^{\pi, \gamma_0}; f_T^{\pi, \gamma_1})}, \end{aligned} \quad (39)$$

where in the second step we used the fact that $\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma_1})$ is non-decreasing in t and $N_T := \sum_{t=1}^T \sum_{l=1}^L n_{lt}$. Previously we established the lower bound (38), which reads as

$$\text{Reg}_T^{\pi, \gamma_0} \geq \frac{C_2}{(\gamma_0 - \gamma)^2} \text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma}).$$

Putting, $\gamma = \gamma_1 = \gamma_0 + 1/4T^{1/4}$ we get $\text{Reg}_T^{\pi, \gamma_0} \geq C_2 \sqrt{T} \text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma_1})$. Combining this bound with (39), we get

$$\begin{aligned} \max_{\gamma \in \{\gamma_0, \gamma_1\}} \text{Reg}_T^{\pi, \gamma} &\geq \frac{1}{2} (\text{Reg}_T^{\pi, \gamma_0} + \text{Reg}_T^{\pi, \gamma_1}) \\ &\geq C \sqrt{T} \left(\text{KL}(f_t^{\pi, \gamma_0}; f_t^{\pi, \gamma_1}) + \frac{N_T}{T} e^{-\text{KL}(f_T^{\pi, \gamma_0}; f_T^{\pi, \gamma_1})} \right) \\ &\geq C \sqrt{T} (1 + \log(N_T/T)), \end{aligned}$$

where in the last step we used the inequality $ae^{-b} + b \geq 1 + \log(a)$. ■

Appendix C. Proofs of all other results and intermediate steps

C.1 Proof of Proposition 1

Recall that $V_{lt}^2 = \|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 \tau^2 + \sigma^2$. Since $W \succeq \mathbf{0}$, and $\rho_t \geq 0$, we have $\mathbf{I} - \rho_t \mathbf{W} \preceq \mathbf{I}$. In addition, by Assumption 2.2, we have $\mathbf{I} - \rho_t \mathbf{W} \succeq \varepsilon \mathbf{I}$. Therefore, since $\|\mathbf{e}_l\| = 1$,

$$1 \leq \|(I - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\| \leq \frac{1}{\varepsilon},$$

from which we obtain the result.

Next, to prove the upper bound on the optimal prices, we recall that

$$0 \leq \mathbf{x}'_{lt} \mathbf{m}_{lt} \leq \|\mathbf{x}_{lt}\| \|\mathbf{m}_{lt}\| \leq \frac{\|\boldsymbol{\mu}_t\|}{V_{lt}} \leq \frac{C_\mu}{c_V}.$$

Invoking relation (14), and noting that β_t and so b_t are negative, we arrive at

$$p_{lt}^* = \frac{1}{-b_{lt}} (\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}) \leq c_\beta^{-1} C_V (C_\mu c_V^{-1} - 0.5\phi(0)),$$

where we used that φ^{-1} is increasing, $\mathbf{x}_{lt} \geq 0$, and $\varphi^{-1}(0) = -0.5/\phi(0)$. ■

C.2 Proof of Lemma 3

By definition, $V_{lt}^2 = \|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 \tau^2 + \sigma^2$. Hence, if ω_* is the smallest eigenvalue of $\mathbf{W}$, then $V_{lt}^2 \geq \tau^2/(1 - \rho_t \omega_*)^2$.

We want to bound the $|b_{l,t+1} - b_{lt}|$ and $\|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2$.

$$\begin{aligned} \|\mathbf{m}_{l,t+1} - \mathbf{m}_{lt}\|_2 &= \left\| \frac{\boldsymbol{\mu}_{t+1}}{V_{l,t+1}} - \frac{\boldsymbol{\mu}_t}{V_{lt}} \right\|_2 \\ &\leq \left\| \frac{\boldsymbol{\mu}_{t+1} - \boldsymbol{\mu}_t}{V_{l,t+1}} \right\|_2 + \boldsymbol{\mu}_t \left\{ \frac{1}{V_{l,t+1}} - \frac{1}{V_{lt}} \right\} \\ &\leq \frac{\delta_{t\mu}}{\tau/(1 - \rho_t \omega_*)} + C_\mu \left\{ \frac{1}{V_{l,t+1}} - \frac{1}{V_{lt}} \right\}. \end{aligned}$$

Further, the second term can be simplified as

$$\begin{aligned} \left\{ \frac{1}{V_{l,t+1}} - \frac{1}{V_{lt}} \right\} &= \frac{V_{lt} - V_{l,t+1}}{V_{l,t+1} V_{lt}} = \frac{V_{lt}^2 - V_{l,t+1}^2}{V_{l,t+1} V_{lt} (V_{lt} + V_{l,t+1})} \\ &\leq \frac{1}{2c_V^3} (V_{lt}^2 - V_{l,t+1}^2) \\ &\leq \frac{\tau^2}{2c_V^3} (\|(\mathbf{I} - \rho_t \mathbf{W})^{-1} \mathbf{e}_l\|^2 - \|(\mathbf{I} - \rho_{t+1} \mathbf{W})^{-1} \mathbf{e}_l\|^2) \leq C \delta_{t\rho}. \end{aligned}$$

The same analysis can be done for $|b_{l,t+1} - b_{lt}|$ as well. ■

C.3 Proof of Corollary 4

The corollary follows directly by applying the results from Lemma 3 in Theorem 2. Since $\rho_t = \rho$ for all t , hence $\delta_{t\rho} = 0$.

Since $\eta_t \propto 1/\sqrt{t}$, we get

- $\mathcal{R}_1 = LC_1 C \tau^{-1} (1 - \rho \omega_*) \sum_{t=1}^T \sqrt{t} \delta_{t\beta},$
- $\mathcal{R}_2 = LC_1 C \tau^{-1} (1 - \rho \omega_*) \sum_{t=1}^T \sqrt{t} \delta_{t\mu},$
- $\mathcal{R}_3 = C_3 C \sum_{t=1}^T n_t^2 / \sqrt{t} = \mathcal{O}(\sqrt{T}),$
- $\mathcal{R}_4 = C_4 CL (C_b + C_m) \sqrt{T+1} = \mathcal{O}(\sqrt{T}).$

Changing the constants appropriately gives us the corollary. ■

C.4 Proof of Lemma 6

Consider the particular set-up when $\mu_t = 0$, $\beta_t = \beta$, $\rho_t = \rho$ and $\sigma = 1$ in (2). Further, assume $n_{lt} = n$ for all l, t and $n \rightarrow \infty$. The proof can easily be extended to the generic set-up. Under these parametric assumptions, first note that, $\text{Rev}(\lambda, l, t, p_{lt}) = np_{lt} \Phi(\alpha_{lt} + \beta p_{lt})$. The optimal pricing strategy p_{lt}^* maximizes $\text{Rev}(\lambda, l, t, p_{lt})$ over p_{lt} for any fixed λ .

Now, consider an arbitrary pricing policy $\mathbf{p}$ based on the unpenalized likelihood $\text{PL}(\lambda, 0)$. Such a policy will be dominated by its oracle counter-part $\mathbf{p}^{\text{or}}$ which already knows the price coefficient β and also, knows the latent utility U_{lt} . Note that, the revenue of any pricing policy $\mathbf{p}$ based on the unpenalized likelihood is always dominated by the revenue of this oracle strategy, i.e., $\text{Rev}(\lambda, \mathbf{p}) \leq \text{Rev}(\lambda, \mathbf{p}^{\text{or}})$. Subsequently, the oracle strategy $\mathbf{p}^{\text{or}}$ will have a lower regret. Next, we concentrate on the regret of $\mathbf{p}^{\text{or}}$.

For this calculation note that based on model (2), the only unknown parameters for the oracle strategy $\mathbf{p}^{\text{or}}$ are the α_{lt} s. Under this framework consider α_{lt} s being best estimated by $\hat{\alpha}_{lt}^{\text{or}}$.

Now, note that as we do not have any structural assumption between α_t and α_{t+1} over $t = 1, \dots, T$, for any t , $\hat{\alpha}_{lt}^{\text{or}}$ will be estimated based on $\{U_{ltk} : l = 1, \dots, L; k = 1, \dots, n\}$. As the prices p_{lt} are known (based on the filtration $\mathcal{F}_{t-1}$ which contains all information upto time $t-1$) this further reduces to estimating the L means α_t from uncorrelated L dimensional Gaussian location model where we observe $n^{-1} \sum_{k=1}^n U_{ltk} - \beta p_{lt}$ for $l = 1, \dots, L$. From the Cramer-Rao lower bound for Gaussian family, it follows that for all $l = 1, \dots, L$, we will have the following error bound on any estimate $\hat{\alpha}_{lt}$:

$$\mathbb{E}_{\lambda} (\hat{\alpha}_{lt} - \alpha_{lt})^2 \geq n^{-1}.$$

As such consider the α_{lt} s under the oracle framework to be estimated by the MLE. Let $\hat{\delta}_{lt} = \hat{\alpha}_{lt}^{\text{or}} - \alpha_{lt}$. Then, noting that the MLE is asymptotically rotation invariant in this case, we have for any λ :

$$\mathbb{E}_{\lambda} \hat{\delta}_t \hat{\delta}_t^T = n^{-1} I_L \quad \text{and} \quad \mathbb{E}_{\lambda} \hat{\delta}_t \rightarrow \mathbf{0} \quad \text{as } n \rightarrow \infty. \quad (40)$$

Now, note that for the oracle strategy,

$$\text{Rev}(\lambda, l, t, p_{lt}^{\text{or}}) = \max_{p \geq 0} np \Phi(\hat{\alpha}_{lt}^{\text{or}} + \beta p) = \max_{p \geq 0} np \Phi(\hat{\delta}_{lt} + \alpha_{lt} + \beta p).$$

Let $f(l, t, p) = np \Phi(\hat{\delta}_{lt} + \alpha_{lt} + \beta p)$. Consider Taylor-Series expansion:

$$f(l, t, p) = np \Phi(\alpha_{lt} + \beta p) + n \hat{\delta}_{lt} p \phi(\alpha_{lt} + \beta p) + 2^{-1} np \hat{\delta}_{lt}^2 \phi'(\alpha_{lt} + \beta p) + r(l, t, p),$$

where $r(l, t, p)$ contains third and higher order terms. Now, we have $L^{-1} \sum_l f(l, t, p_{lt})$ converges in probability to

$$\frac{1}{L} \sum_l np_{lt} \Phi(\alpha_{lt} + \beta p_{lt}) + \frac{1}{L} \sum_l n p_{lt} \phi(\alpha_{lt} + \beta p_{lt}) \mathbb{E}_{\lambda} \hat{\delta}_{lt} + \frac{1}{2L} \sum_l np_{lt} \phi'(\alpha_{lt} + \beta p_{lt}) \mathbb{E}_{\lambda} \hat{\delta}_{lt}^2,$$

as $L^{-1} \sum_l r(l, t, p_{lt}) \rightarrow 0$ in probability as $n \mathbb{E}_{\lambda} \hat{\delta}_{lt}^{2+m} = O(n^{-m/2})$, for $m \geq 1$. Using (40), the second term in the above expression vanishes and the third term gets further simplified, resulting in the following asymptotic result:

$$\frac{1}{L} \sum_l f(l, t, p_{lt}) = n \left[\frac{1}{L} \sum_l p_{lt} \Phi(\alpha_{lt} + \beta p_{lt}) \right] + \frac{1}{2L} \sum_l p_{lt} \phi'(\alpha_{lt} + \beta p_{lt}) + o(1).$$

Thus, the regret of $\mathbf{p}^{\text{or}}$ at time t is given by

$$L^{-1} \sum_{l=1}^L \mathcal{R}_{lt}(\lambda, \mathbf{p}^{\text{or}}) \geq (\mathcal{A} - \mathcal{B})/L + o(1), \quad (41)$$

where,

$$\begin{aligned} \mathcal{A} &= \max_{p_{lt}: l=1, \dots, L} \left[\sum_l np_{lt} \Phi(\alpha_{lt} + \beta p_{lt}) \right], \text{ and} \\ \mathcal{B} &= \max_{p_{lt}: l=1, \dots, L} \left[\sum_l np_{lt} \Phi(\alpha_{lt} + \beta p_{lt}) - 2^{-1} \sum_l p_{lt} (\alpha_{lt} + \beta p_{lt}) \phi(\alpha_{lt} + \beta p_{lt}) \right]. \end{aligned}$$

Note that, the expression in $\mathcal{B}$ is simplified using $\phi'(u) = -u\phi(u)$. Now recall that β , being the price sensitivity, is negative. Based on model (2), for the utilities to be positive we have the following assumption of the price: $\alpha_{lt} + \beta p_{lt} > 0$ for all l and t . Let the prices be selected such that $\inf_l \alpha_{lt} + \beta p_{lt} > \epsilon_0$ for some prefixed small $\epsilon_0 > 0$. By Proposition 1, the optimal prices are bounded and so are $\sup_l \alpha_{lt} + \beta p_{lt} < M_0$. Then,

$$\mathcal{A} - \mathcal{B} \geq 2^{-1} \epsilon \sum_l p_{lt}^*,$$

where p_{lt}^* is the optimal price based on criterion $\mathcal{B}$, and $\epsilon = \min_{\epsilon_0 < |u| < M'} u\phi(u)$. Thus, the cumulative regret of $\mathbf{p}^{\text{or}}$ over time is given by

$$\mathcal{R}(\lambda, \mathbf{p}^{\text{or}}) = \sum_{t=1}^T \sum_{l=1}^L \mathcal{R}_{lt}(\lambda, \mathbf{p}^{\text{or}}) = \Omega(LT).$$

Thus, we have,

$$\mathcal{B}(\theta, \mathbf{p}_U) = \Omega(LT).$$

Now, consider the regret from the proposed strategy. Based on (24), we have

$$\mathcal{B}(\boldsymbol{\theta}, \mathbf{p}) \leq C_6 \tau^{-1} (1 - \rho_* \omega_*) \sum_{t=1}^T \sqrt{t} (\delta_{t\beta} + \delta_{t\mu}) + C_7 \sum_{t=1}^T \sqrt{t} \delta_{t\rho} + \mathcal{O}(\sqrt{T}) = \mathcal{O}(\sqrt{T}),$$

where, the second asymptotic result follows as $\sum_{t=1}^T \sqrt{t} \delta_{t\beta}$, $\sum_{t=1}^T \sqrt{t} \delta_{t\mu}$ and $\sum_{t=1}^T \sqrt{t} \delta_{t\rho}$ are all bounded above by $\mathcal{O}(\sqrt{T})$. Comparing the above two displays the result follows. $\blacksquare$

C.5 Proof of Proposition 7

Consider the revenue function given by (25), $\text{Rev}_{lt}(p) = n_{lt} p \Phi(b_{lt} p + \mathbf{x}'_{lt} \mathbf{m}_{lt})$. By definition, p_{lt}^* is the maximizer of $\text{Rev}_{lt}(p)$ and hence $\text{Rev}'_{lt}(p_{lt}^*) = 0$. Using Taylor series expansion around p_{lt}^* , we get

$$\mathcal{R}_{lt} = \text{Rev}_{lt}(p_{lt}^*) - \text{Rev}_{lt}(p_{lt}) = \frac{1}{2} \text{Rev}''_{lt}(p)(p_{lt} - p_{lt}^*)^2,$$

for some p between p_{lt} and p_{lt}^* . The second derivative can be bounded as

$$\frac{1}{2} \text{Rev}''_{lt}(p) = n_{lt} \frac{2b_{lt}\phi(b_{lt}p + \mathbf{x}'_{lt} \mathbf{m}_{lt}) + pb_{lt}^2 \phi'(b_{lt}p + \mathbf{x}'_{lt} \mathbf{m}_{lt})}{2} \leq \left(C_b \phi(0) + MC_b^2 \frac{\phi(0)}{\sqrt{2}} \right) n_{lt}.$$

Setting $C_9 = C_b \phi(0) + MC_b^2 \frac{\phi(0)}{\sqrt{2}}$ completes the proof. $\blacksquare$

C.6 Proof of Lemma 8

We have the true parameters b_{lt} , $\mathbf{m}_{lt}$ and the output $\hat{b}_{lt}$, $\hat{\mathbf{m}}_{lt}$ from our PSGD pricing policy. Also, p_{lt} and p_{lt}^* are the price based on our policy and the optimal price based on the true parameters, respectively.

As discussed in section 2, we can write the prices in terms of the utility model parameters using the function $g(\cdot, \cdot)$, as follows:

$$\begin{aligned} p_{lt}^* &:= g(b_{lt}, \mathbf{m}_{lt}) = -\frac{\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}}{b_{lt}}, \\ p_{lt} &:= g(\hat{b}_{lt}, \hat{\mathbf{m}}_{lt}) = -\frac{\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}}{\hat{b}_{lt}}. \end{aligned}$$

Now, note that,

$$(p_{lt} - p_{lt}^*)^2 = \left\{ (\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}) b_{lt}^{-1} - (\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) \hat{b}_{lt}^{-1} \right\}^2 = \{A + B\}^2,$$

where, the right side above is decomposed as

$$\begin{aligned} A &= (\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}) b_{lt}^{-1} - (\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) b_{lt}^{-1}, \text{ and,} \\ B &= (\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) (1/b_{lt} - 1/\hat{b}_{lt}). \end{aligned}$$

Using the naive bound $\{A + B\}^2 \leq 2(A^2 + B^2)$ first and then the bound $|b_{lt}| = |\beta_t|/V_{lt} \geq c_\beta/C_V$ and the policy rule $p_{lt} = g(\hat{b}_{lt}, \hat{\mathbf{m}}_{lt})$, it follows that $(p_{lt} - p_{lt}^*)^2$ is bounded above by

$$2C_V^2 c_\beta^{-2} \{(\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}) - (\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt})\}^2 + 2p_{lt}^2 b_{lt}^{-2} (b_{lt} - \hat{b}_{lt})^2,$$

Since, $\varphi^{-1}(-v) + v$ is 1-Lipschitz, we have

$$((\varphi^{-1}(-\mathbf{x}'_{lt} \mathbf{m}_{lt}) + \mathbf{x}'_{lt} \mathbf{m}_{lt}) - (\varphi^{-1}(-\mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}) + \mathbf{x}'_{lt} \hat{\mathbf{m}}_{lt}))^2 \leq \langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2.$$

Therefore, we have

$$(p_{lt} - p_{lt}^*)^2 \leq 2C_V^2 c_\beta^{-2} \langle \mathbf{x}_{lt}, \mathbf{m}_{lt} - \hat{\mathbf{m}}_{lt} \rangle^2 + 2C_V^2 c_\beta^{-2} p_{lt}^2 (b_{lt} - \hat{b}_{lt})^2.$$

Setting $C_{10} = 2C_V^2 c_\beta^{-2}$ proves the lemma. ■

References

- L. Anselin. *Spatial econometrics: methods and models*, volume 4. Springer Science & Business Media, 2013.
- G.-Y. Ban and N. B. Keskin. Personalized dynamic pricing with machine learning: High-dimensional features and heterogeneous elasticity. *Management Science*, 67(9):5549–5568, 2021.
- S. Banerjee, B. P. Carlin, and A. E. Gelfand. *Hierarchical modeling and analysis for spatial data*. CRC press, 2014.
- T. Banerjee, G. Mukherjee, and D. Paul. Improved shrinkage prediction under a spiked covariance structure. *J. Mach. Learn. Res.*, 22:180–1, 2021.
- T. Banerjee, P. Liu, G. Mukherjee, S. Dutta, and H. Che. Joint modeling of playing time and purchase propensity in massively multiplayer online role playing games using crossed random effects. *Annals of Applied Statistics*, 1(1):1–22, 2022.
- B. Berman and S. Thelen. Planning and implementing an effective omnichannel marketing program. *International Journal of Retail & Distribution Management*, 46(7):598–614, 2018.
- O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.
- O. Besbes, Y. Gur, and A. Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- K. Bimpikis, O. Candogan, and D. Saban. Spatial pricing in ride-sharing networks. *Operations Research*, 67(3):744–769, 2019.
- E. T. Bradlow, B. Bronnenberg, G. J. Russell, N. Arora, D. R. Bell, S. D. Duvvuri, F. T. Hofstede, C. Sismeiro, R. Thomadsen, and S. Yang. Spatial models in marketing. *Marketing letters*, 16:267–278, 2005.
- Y. Bramoullé, H. Djebbari, and B. Fortin. Identification of peer effects through social networks. *Journal of econometrics*, 150(1):41–55, 2009.
- J. Broder and P. Rusmevichientong. Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980, 2012.
- L. D. Brown, G. Mukherjee, and A. Weinstein. Empirical bayes estimates for a two-way cross-classified model. *The Annals of Statistics*, 46(4):1693–1720, 2018.
- C. Bureau. *Statistical Abstract of the United States 2008*. Government Printing Office, 2008.
- N. Cesa-Bianchi, C. Gentile, and Y. Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory*, 61(1):549–564, 2015.
- W. C. Cheung, D. Simchi-Levi, and H. Wang. Dynamic pricing and demand learning with limited price experimentation. *Operations Research*, 65(6):1722–1731, 2017.

- P. Chintagunta, J.-P. Dube, and V. Singh. Market structure across stores: An application of a random coefficients logit model with store level data. In *Advances in Econometrics*. Emerald Group Publishing Limited, 2002.
- H. Choi, C. F. Mela, S. R. Balseiro, and A. Leary. Online display advertising markets: A literature review and future directions. *Information Systems Research*, 31(2):556–575, 2020.
- M. C. Cohen, I. Lobel, and R. Paes Leme. Feature-based dynamic pricing. *Management Science*, 66(11):4921–4943, 2020.
- T. M. Cover and J. A. Thomas. Information theory and statistics. *Elements of information theory*, 1(1):279–335, 1991.
- T. H. Cui, A. Ghose, H. Halaburda, R. Iyengar, K. Pauwels, S. Sriram, C. Tucker, and S. Venkataraman. Informational challenges in omnichannel marketing: remedies and future research. *Journal of Marketing*, 85(1):103–120, 2021.
- A. V. den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015.
- A. V. den Boer and B. Zwart. Simultaneously learning and optimizing using controlled variance pricing. *Management science*, 60(3):770–783, 2013.
- B. Efron and T. Hastie. *Computer Age Statistical Inference: Algorithms, Evidence, and Data Science*, volume 6. Cambridge University Press, 2021.
- V. F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Operations Research*, 58(1):16–29, 2010.
- K. J. Ferreira, D. Simchi-Levi, and H. Wang. Online network revenue management using thompson sampling. *Operations research*, 66(6):1586–1602, 2018.
- A. S. Fotheringham, C. Brunson, and M. Charlton. *Geographically weighted regression: the analysis of spatially varying relationships*. John Wiley & Sons, 2003.
- D. Fourdrinier, W. E. Strawderman, and M. T. Wells. *Shrinkage estimation*. Springer, 2018.
- A. E. Gelfand, P. Diggle, P. Guttorp, and M. Fuentes. *Handbook of spatial statistics*. CRC press, 2010.
- N. Golrezaei, H. Nazerzadeh, and R. Randhawa. Dynamic pricing for heterogeneous time-sensitive customers. *Manufacturing & Service Operations Management*, 22(3):562–581, 2020.
- N. Golrezaei, A. Javanmard, and V. S. Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Oper. Res.*, 69(1):297–314, 2021.
- J. M. Harrison, N. B. Keskin, and A. Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. *Management Science*, 58(3):570–586, 2012.

- T. Hastie, R. Tibshirani, J. H. Friedman, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.
- C.-S. Hsieh and L. F. Lee. A social interactions model with endogenous friendship formation and selectivity. *Journal of Applied Econometrics*, 31(2):301–319, 2016.
- A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, pages 398–406. PMLR, 2015.
- W. Jank and P. Kannan. Understanding geographical markets of online firms using spatial models of customer choice. *Marketing Science*, 24(4):623–634, 2005.
- A. Javanmard. Perishability of data: dynamic pricing under varying-coefficient models. *The Journal of Machine Learning Research*, 18(1):1714–1744, 2017.
- A. Javanmard and H. Nazerzadeh. Dynamic pricing in high-dimensions. *The Journal of Machine Learning Research*, 20(1):315–363, 2019.
- B. Karmakar, O. Kwon, G. Mukherjee, S. Siddarth, and J. M. Silva-Risso. Does a consumer’s previous purchase predict other consumers’ choices? a bayesian probit model with spatial correlation in preference. *under review in QME; available at: bit.ly/spatialprobit*, 2021.
- N. B. Keskin and A. Zeevi. Chasing demand: Learning and earning in a changing environment. *Mathematics of Operations Research*, 42(2):277–307, 2016.
- N. B. Keskin and A. Zeevi. Chasing demand: Learning and earning in a changing environment. *Mathematics of Operations Research*, 42(2):277–307, 2017.
- S. Kou and J. J. Yang. Optimal shrinkage estimation in heteroscedastic hierarchical linear models. In *Big and Complex Data Analysis*, pages 249–284. Springer, 2017.
- R. P. Leme and J. Schneider. Contextual search via intrinsic volumes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 268–282. IEEE, 2018.
- J. P. LeSage. A family of geographically weighted regression models. In *Advances in spatial econometrics*, pages 241–264. Springer, 2004.
- Y. Liu-Thompkins. A decade of online advertising research: What we learned and what we need to know. *Journal of advertising*, 48(1):1–13, 2019.
- I. Lobel, R. P. Leme, and A. Vladu. Multidimensional binary search for contextual decision-making. *arXiv preprint [arXiv:1611.00829](https://arxiv.org/abs/1611.00829)*, 2016.
- L. Ma, R. Krishnan, and A. L. Montgomery. Latent homophily or social influence? an empirical analysis of purchase within a social network. *Management Science*, 61(2):454–473, 2015.

- C. F. Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.
- S. Mukhopadhyay, W. Kar, and G. Mukherjee. Estimating promotion effects in email marketing using a large-scale cross-classified bayesian joint model for nested imbalanced data. *Annals of Applied Statistics*, 1(1):1–21, 2022.
- R. B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- J. P. Peter, J. H. Donnelly, and M. B. Vandenbosch. *A preface to marketing management*. McGraw-Hill, 2000.
- E. M. Schwartz, E. T. Bradlow, and P. S. Fader. Customer acquisition via display advertising using multi-armed bandit experiments. *Marketing Science*, 36(4):500–522, 2017.
- Z. Tan. Improved minimax estimation of a multivariate normal mean under heteroscedasticity. *Bernoulli*, 21(1):574–603, 2015.
- K. E. Train. *Discrete choice methods with simulation*. Cambridge university press, 2009.
- A. B. Tsybakov. Introduction to nonparametric estimation, 2009. URL <https://doi.org/10.1007/b13794>. Revised and extended from the, 9(10), 2004.
- A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008. ISBN 0387790519, 9780387790510.
- A. Weinstein, Z. Ma, L. D. Brown, and C.-H. Zhang. Group-linear empirical bayes estimates for a heteroscedastic normal mean. *Journal of the American Statistical Association*, 113(522):698–710, 2018.
- X. Xie, S. Kou, and L. D. Brown. Sure estimates for a heteroscedastic hierarchical model. *Journal of the American Statistical Association*, 107(500):1465–1479, 2012.
- S. Yang and G. M. Allenby. Modeling interdependent consumer preferences. *Journal of Marketing Research*, 40(3):282–294, 2003.
- T. Yang, L. Zhang, R. Jin, and J. Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *International Conference on Machine Learning*, pages 449–457. PMLR, 2016.
- J. Zhou, Y. Tu, Y. Chen, and H. Wang. Estimating spatial autocorrelation with sampled network data. *Journal of Business & Economic Statistics*, 35(1):130–138, 2017.
- Z. Zhou, R. Xu, and J. Blanchet. Learning in generalized linear contextual bandits with stochastic delays. *Advances in Neural Information Processing Systems*, 32, 2019.
- M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pages 928–936, 2003.