MICHIGAN ROSS

Working Paper

Dynamic Joint Pricing and Order Fulfillment for E-Commerce Retailers

Yanzhe Lei
Stephen M. Ross School of Business
University of Michigan

Stefanus Jasin
Stephen M. Ross School of Business
University of Michigan

Amitabh Sinha
Stephen M. Ross School of Business
University of Michigan

Ross School of Business Working Paper
Working Paper No. 1310
March 2016

Dynamic Joint Pricing and Order Fulfillment for E-commerce Retailers

Yanzhe (Murray) Lei, Stefanus Jasin, Amitabh Sinha
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI 48109,
leiyz, sjasin, amitabh@umich.edu

We consider an e-commerce retailer (e-tailer) who sells a catalog of products to customers from different regions during a finite selling season and fulfills orders through multiple fulfillment centers. The e-tailer faces a *Joint Pricing and Fulfillment* (JPF) problem: At the beginning of each period, she needs to jointly decide the price for each product and how to fulfill an incoming order. The objective is to maximize the total expected profits defined as total expected revenues minus total expected shipping costs (all other costs are fixed in this problem). The exact optimal policy for JPF is difficult to solve; so, we propose two heuristics that have provably good performance compared to reasonable benchmarks. Our first heuristic directly uses the solution of a deterministic approximation of JPF as its control parameters whereas our second heuristic improves the first heuristic by adaptively adjusting the original control parameters at the beginning of every period. An important feature of the second heuristic is that it decouples the pricing and fulfillment decisions, making it easy to implement. We show theoretically and numerically that the second heuristic significantly outperforms the first heuristic and is very close to a benchmark that jointly re-optimizes the full deterministic problem at every period.

Key words: dynamic pricing, fulfillment policies, e-commerce retail, asymptotic analysis.

1. Introduction

Driven by the growing population of internet users, the retailing industry has witnessed a boom in the e-commerce channel during the past decades. According to U.S. Census Bureau (2016), for the year of 2015, the sales of e-commerce retail in the United States grows continually at an impressive rate of 14.63%, which accounted for 68% of the growth of the U.S. retail sector. While the growth statistics are impressive, it does not mean that online retailing is an easy business to run. As pointed out in Rigby (2014), Amazon.com, whose figure is similar to other e-tailers, has averaged only 1.3% in operating margin over the past three years; in contrast, the operating margin for department/discount stores typically run about 6% to 10%. Despite its razor-thin margin, e-tailers have to spend heavily in expenditure to meet consumers' evolving expectation. For example, building a highly productive fulfillment center (FC), the typical facility through which e-tailers handle in/outbound logistic, costs at least \$250 million, ten times more than building a large department store (Rigby 2014). All these factors put together highlight the importance for e-tailers to operate in a way that maximizes their revenue while at the same time minimizing their expenditure.

Compared to its brick-and-mortar counterpart, an e-tailer has extra flexibilities in responding to the market by being able to change prices frequently in real-time (Chen 2014) and reduce outbound shipping cost through tactical order fulfillment (Agatz et al. 2008). Indeed, powered by a vast amount of data and efficient IT infrastructure, e-tailers nowadays actively adjust their prices according to the imbalance between supply and demand, and other external factors in the market. This practice, also known as *dynamic pricing*, has been widely adopted in many industries including airlines, car rental, hotel, and cruise. The retailing industry is among the latest incursions, pioneered by Amazon.com, who is reported to adjust its price lists every ten minutes on average (Shpanya 2014). As also reported in the same article, at least 22% of retailers, including Sears, Bestbuy, and Walmart, have chosen to implement automatic pricing solutions in their online channel and improve their gross margin by 10%.

Unlike pricing decisions that are executed online and have an immediate impact on the revenue stream, an e-tailer's fulfillment decisions affect the physical distribution of inventories and have an immediate impact on its operating cost. Among the different parts of an e-tailer's fulfillment plan, outbound shipping is often cited as the primary source of cost (Dinlersoz and Li 2006). For example, Amazon.com spent \$11.54 billion in the fiscal year of 2015 on outbound shipping alone (including sortation and delivery center costs); this roughly represents 10% of its net revenue (\$107.01 billion) and 30% incremental over the total costs in 2014 (\$8.71 billion) (Amazon.com 2015). While consumers value a good fulfillment model, they often do not want to share the cost by paying additional shipping fees. According to Sides and Hogan (2015), 72% of the consumers surveyed cite free shipping as the offering they would take advantage of when shopping online, and 87% of them rank free shipping as being more important than fast shipping. Moreover, an extra charge on delivery can negatively impact consumer's purchase intention. For example, UPS (2014) attributes 50% of shopping cart abandonments to unexpectedly high shipping fee. To mitigate the adverse effects of charging shipping fee, many retailers now offer appealing shipping options for online shoppers such as unconditional free shipping (Nordstrom, Zappos), contingent free shipping (Amazon.com, Jet.com), and free in-store pickup (Macy's, Walmart). As a consequence, e-tailers are strongly incentivized to find the cheapest fulfillment plan on every single order.

Conceptually, pricing and fulfillment decisions are closely tied together, since they both immediately affect the balance between supply and demand. On the one hand, an e-tailer's fulfillment strategy affects her pricing decision as the price that maximizes total revenues does not necessarily maximize total expected profits (i.e., revenue minus cost); on the other hand, the effectiveness of a fulfilment strategy heavily depends on the current inventory distribution and forecasted future demands, which in turn are determined by the pricing decision. This interdependency calls for a systematic study of joint pricing and fulfillment optimization. To illustrate the potential benefit of managing pricing and fulfillment jointly instead of separately in an e-commerce environment, we describe the following simple example. Consider an e-tailer selling a cast-iron grill pan weighing 7.1 lbs to Midwest and West Coast regions. Customers from both regions see the same price posted online. The demand for the grill pan is divisible and deterministic. For the purpose of illustration, we assume a demand function $\lambda(p) = 58 - p$ for both regions. The price is restricted to within the range of \$14.22 and \$30.34 (see Camelcamelcamel.com 2016 for a price history of a similar product at Amazon.com). The e-tailer has a distribution network consisting of two FCs located at California (CA) and Illinois (IL), which hold C_{CA} and C_{IL} unit of inventory, respectively. Each customer purchases exactly one grill pan, which is to be shipped immediately from either FC using UPS' 3-day select service. Figure 1 describes the basic setting of the profit maximization problem faced by the e-tailer, where we use MI (Michigan) and OR (Oregon) as representatives of Midwest and Westcoast regions, respectively. Shipping cost data is gathered from UPS (2016).

Suppose that $C_{IL} = 30$ and $C_{CA} = 28$, i.e., the inventory level in IL is slightly higher than the inventory level in CA. If the e-tailer manages the pricing decision separately from fulfillment assignment (i.e., in subsequent manner), she would first solve a revenue maximization problem: $\max_{p\in[\$14.22,\$30.34]} \{p(58-p)+p(58-p) : (58-p)+(58-p) \leq 10+48\}$. The optimal solution is given by $p = \$29.00$, which results in 29 units of demand from each MI and OR and yields a total revenue of $\$29 \times 29 \times 2 = \$1,682.00$. Next, she needs to decide how to fulfill these orders by solving the following cost minimization problem: $\min_{x_{ij}\geq 0}\{\sum_{i\in\{CA,IL\}}\sum_{j\in\{MI,OR\}} c_{ij}x_{ij} \ : \sum_{i\in\{CA,IL\}} x_{ij} = 29, \forall j, \sum_{j\in\{MI,OR\}} x_{ij} \leq C_i, \forall i\}$. The optimal solution is $x_{IL,MI} = 29, x_{IL,OR} = 1, x_{CA,MI} = 0, x_{CA,OR} = 28$. This yields a total shipping costs of $\$1,123.05$, leaving the e-tailer with a net profit of $\$1,682.00 - \$1,123.05 = \$558.95$. Suppose now that the e-tailer manages the pricing and fulfillment decisions jointly by solving the following profit maximization problem:

$$\begin{aligned}\max_{p\in[\$14.22,\$30.34],\ x_{ij}\geq 0} \quad & p(58-p) + p(58-p) - \sum_{i\in\{CA,IL\}}\sum_{j\in\{MI,OR\}} c_{ij}x_{ij} \\ \text{s.t.} \quad & \sum_{i\in\{CA,IL\}} x_{ij} = 58 - p,\ \forall j, \quad \sum_{j\in\{MI,OR\}} x_{ij} \leq C_i,\ \forall i.\end{aligned}$$

The optimal solution is $p = \$30.34, x_{IL,MI} = x_{CA,OR} = 27.66, x_{IL,OR} = x_{CA,MI} = 0$. The corresponding net profit is $\$624.56$, which is 8.22% larger than the net profit of managing the pricing and fulfillment decisions separately. This is because, although the increment in price lowers the revenue, it also reduces the demand so that we no longer ship on the IL-OR and CA-MI routes which have negative profit margins. It is not difficult to check that increasing inventory imbalance across the two FCs will result in a larger marginal improvement. For example, if we set $C_{IL} = 9$ and $C_{CA} = 49$, the marginal improvement of net profit due to optimizing price and fulfillment assignment jointly is as large as 101.20%.

The above example shows that optimizing pricing and fulfillment decisions jointly could be very effective even when future demands are known exactly. It is safe to conjecture that this benefit will even be larger when the e-tailer is facing uncertain demand. Indeed, even if the initial inventory levels are properly chosen by taking into account both the pricing and

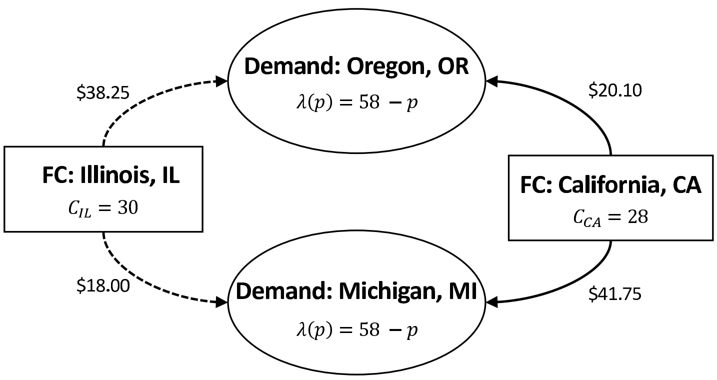

Figure 1 A 2-FC 2-Demand-Location Example.

fulfillment decisions, inventory imbalance is bound to happen within the replenishment cycle due to demand randomness. Moreover, demand for an item may depend not only on its own price, but also on the price of other products that may be complements or substitutes. This suggests that, in order to get the most profit, e-tailers need to continuously check-and-correct for the imbalance across their supply networks, which brings us to the research question of this paper: *How should an e-tailer manage the pricing and fulfillment decision for multiple products jointly by utilizing the information regarding the current inventory distribution and future demand projection in a way that maximizes total expected profits?*

Our results and contributions. We consider a multi-period *Joint Pricing and Fulfillment* (JPF) problem where an e-tailer sells multiple products to customers coming from different demand locations and demands are fulfilled through multiple FCs. The decision variables are the price and fulfillment assignment; the objective is to maximize total expected profits. Our results and contributions in this paper are summarized below:

1. To the best of our knowledge, we are the first to consider the JPF problem. This is surprising given the importance of pricing and fulfillment as tactical levers to maximize total expected profits in e-tail setting. See Section 2 for extensive literature on these individual problems.
2. A distinct feature of e-commerce retail is that the e-tailer cannot price-differentiate customers from different demand locations by charging different prices for the same product during the same period. This constraint introduces complexities that do not previously appear in the relevant literature (see discussions in Section 3). To overcome this problem, we propose a novel deterministic relaxation of the original stochastic control problem where all the random variables are approximated by their expected values and the pricing decision is approximated by a *randomization* over a fixed set of discrete prices. We show that there exists a set of discrete prices such that the optimal value of the resulting *Approximate Linear Program* (ALP) well approximates that of JPF.
3. We first propose a simple heuristic, which we call *Randomized Pricing and Control* (RPF). RPF uses the ALP solution as probabilities to set pricing and fulfillment decisions at each time period. Although this is a *static* heuristic (it uses the same probabilities throughout the selling season), we show in our numerical studies in Section

7 that it dominates a benchmark policy that determines the pricing and fulfillment decisions separately, even if those are re-optimized at the beginning of each period.

4. Next, we refine RPF by adaptively adjusting the pricing and fulfillment decisions in every period: prices are adjusted using a linear control without any re-optimization, while the fulfillment problem is solved as a simple transportation problem that is separable over the products (and thus can be solved in parallel). We prove that this new heuristic (called R^2PF, for *Re-adjust and Re-optimize Pricing and Fulfillment*) has a significantly better performance than RPF; our numerical studies confirm and quantify this.
5. Methodologically, our work contributes to the literature by (1) proposing a deterministic approximation of JPF problem, (2) generalizing existing works on one-point adjustment to distribution adjustment (see Section 2 for more discussions), and (3) proposing a novel combination of a real-time adjustment of some decision variables with a re-optimized update of other decision variables. We think that this combination of separately adjusting two subsets of decision variables can also be useful for other applications where the number of decision variables is large and the problem has some structure that can be exploited to do this.
6. Managerially, our work offers an interesting insight: It highlights the potential benefit of an effective top-down policy for managing both demand (via pricing) and supply (via fulfillment). The purpose of the first stage of R^2PF is to maintain balance between supply and demand at an aggregate level, between total available inventories at all FCs and total forecasted future demands from all locations. The second stage of R^2PF deals with what is left of the first stage: It takes into account the actual inventory distribution across different FCs and computes a fulfillment assignment that minimizes total shipping costs. These two stages are, in general, indispensable. Without the aggregate re-balancing in the first stage, the fulfillment optimization in the second stage will only be minimizing shipping cost without maximizing revenue; without the fulfillment optimization in the second stage, the aggregate re-balancing in the first stage may result in a high shipping cost, which leads to a lower net profit.

Organization of the paper. The related literature is reviewed in Section 2. In Section 3, we formally formulate the JPF problem and state our modeling assumptions. We propose an approximation scheme and our performance measure in Section 4. Sections 5 and 6 are devoted to the analysis of our heuristics. Numerical experiments are presented in Section 7. Finally, in Section 8, we conclude the paper. All the proof of the results and the parameters of the numerical experiments can be found in the electronic companion for this paper.

2. Literature Review

In terms of topic, the problem that we study in this paper is related to three streams of literature: dynamic pricing, e-commerce fulfillment policies, and the interaction between pricing and fulfillment-related decisions. In terms of methodology, our work is related to the study of asymptotic performance of re-optimization-based heuristic and linear control (or real-time adjustment) heuristic. We discuss them in turn.

Dynamic Pricing. In the revenue management (RM) literature, research on dynamic pricing studies how a firm should dynamically change their price to balance supply and demand during a finite selling season; see Talluri and van Ryzin (2006) and Özer and Phillips (2012) for comprehensive reviews. Although the idea was popularized by its application in airline ticket pricing, as argued by Boyd and Bilegan (2003), the classic dynamic pricing model can also cover the revenue maximization problem in e-commerce. Several works discuss how to design an optimal pricing policy for specific types of e-tailer's problems. For example, Netessine et al. (2006) and Aydin and Ziya (2008) explore the optimal policy for dynamic pricing and packaging when an e-tailer offers an additional product other than the product requested by consumers as a bundle; Ferreira et al. (2015) and Fisher et al. (2015) devise pricing decision support systems for large e-tailers and illustrate their effectiveness by conducting field experiments. Compared to the existing models in the RM literature and the papers cited above, our model shares similarity in the price-induced nature of demand generation and some related assumptions (see Section 3). Unlike the existing literature, though, we jointly consider both the pricing and fulfillment decisions.

E-commerce Fulfillment Policies. The advent of e-commerce has led to substantial research in various aspects of optimizing e-commerce supply chains; see Simchi-Levi et al. (2004) and Agatz et al. (2008) for comprehensive reviews. The fulfillment part of our model focuses exclusively on designing an outbound shipping assignment strategy that helps the e-tailer minimize total shipping costs. A similar problem was first studied by Xu et al. (2009); they construct a heuristic that periodically re-evaluates the real-time assignment decisions based on the currently available information, and illustrate its effectiveness using numerical experiments. Their objective is to minimize the number of split shipments. Acimovic and Graves (2014) study a similar problem and develop a heuristic that minimizes total shipping costs instead of the number of split shipments. Using industry data, they show that their heuristic captures 36% of the savings on costs induced by the optimal hindsight heuristic. More recently, Jasin and Sinha (2015) consider a multi-item fulfillment cost minimization problem. They first propose a heuristic based on the solution of a deterministic relaxation linear program (LP). They then show how to improve the performance of the first heuristic by carefully constructing a correlated rounding scheme and prove its theoretical performance guarantee. Since our focus in this work is on the benefit of joint optimization of pricing and fulfillment decisions, for the fulfillment part, we simplify the model in Jasin and Sinha (2015) by requiring that each order consists of exactly one item. However, the additional layer of the pricing decision, as well as the re-adjusting/re-optimization feature of our main heuristic, precludes a direct generalization of the methodology used in Jasin and Sinha (2015).

Interaction between pricing and fulfillment-related decisions. There have been a few works that study the interplay between e-tailer's pricing decisions and shipping policy, i.e., the format and the extra fee charged on deliveries. Leng and Becerril-Arreola (2010) investigate the impact of contingent free-shipping policy on consumers' purchase decision and derived optimal static pricing and the free-shipping cut-off for e-tailers; Becerril-Arreola et al. (2013) extend the model analyzed in Leng and Becerril-Arreola (2010) by incorporating a second-stage inventory level decision and study the problem by a simulation-based analysis. Gümüş et al. (2013) develop a game-theoretic model to study whether it is optimal for the e-tailer to charge a separate shipping charge, or to incorporate it in the product price but offer a free-shipping policy. In our work, we do not explicitly consider the issue of designing a shipping policy (the format and the extra charge for deliveries); instead, we simply assume a certain cost structure and analyze how to dynamically adjust both the price and fulfillment decisions given the structure.

Re-optimization-based heuristics. In the broader dynamic optimization literature where a multi-period stochastic control problem is often intractable, re-optimization is typically used as a heuristic approach due to its simplicity. Roughly speaking, a re-optimization-based heuristic first approximates the original stochastic control problem with a simple optimization problem (e.g., an LP); as time evolves and uncertainties are realized, the heuristic re-optimizes the approximate optimization problem by updating its parameters to the status quo. In the operations management (OM) literature, this idea has been applied to price-based RM (Maglaras and Meissner 2006, Jasin 2014), quantity-based RM (Reiman and Wang 2008, Ciocan and Farias 2012, Jasin and Kumar 2012, 2013), inventory control (Plambeck and Ward 2006, Secomandi 2008, Doğru et al. 2010, Ahn et al. 2015), and vehicle routing (Secomandi and Margot 2009). Our main heuristic shares the same spirit with existing literature. However, there are some subtleties that differentiate previous works from ours. First, in our setting, it is not trivial to construct a proper deterministic relaxation that both well approximates the original stochastic control and motivates a practical heuristic (see Section 4). We propose a sequence of approximation schemes that trade-offs computational complexity with approximation quality. Second, since our approximate optimization can be very large in size for a high-quality approximation, frequent re-optimizations may not be practically feasible. Thus, we introduce a new methodological novelty by decoupling the pricing and fulfillment decisions. For our main heuristic, only the fulfillment assignment decisions involve re-solving an LP. The size of this LP is much smaller than the original approximate optimization problem and is decomposable over the products. This makes the re-optimization part of our heuristic very time-efficient.

Linear control (real-time adjustment) heuristics. Broadly speaking, a linear control prescribes that the current decision rule can be calculated as an affine function of a baseline control and realized historical outcomes. Similar to re-optimization-based heuristics, linear control is often used as a heuristic approach to deal with an intractable multi-period stochastic control problem. It also has been widely applied in different applications including robust optimization (Ben-Tal et al. 2004, Bertsimas et al. 2010), portfolio management (Calafiore 2009, Moallemi and Saglam 2012), and dynamic pricing (Atar and Reiman 2012, Jasin 2014, Chen et al. 2015). Since a linear control typically runs in real-time, without re-solving any optimization, it is very time-efficient and is sometimes preferable to re-optimization-based heuristics. The exact value of the parameters used in a linear control can either be optimized off-line or computed in a specific way to achieve a certain objective. In our main heuristic, the pricing decisions are adjusted according to an autonomous price update scheme akin to the one used in Jasin (2014) and Chen et al. (2015) (see Section 6). (Although our update rule is not exactly linear, it shares the same spirit of real-time adjustment.) However, there is an important difference: In both Jasin (2014) and Chen et al. (2015), the adjustment is made directly to the price of each product whereas, in ours, the adjustment is made to the set of discrete prices from which the actual price will be sampled. Thus, we are essentially generalizing the one-point adjustment scheme in existing literature to a distribution adjustment scheme.

3. Problem Formulation

Consider a monopolistic e-tailer selling a catalog of K products to customers in J locations with sales fulfilled from I FCs. Throughout the paper, we will use $[N]$ to denote the set $\{1, \ldots, N\}$ for any $N \in \mathbb{N}_+$. The selling season is finite and divided into $T \geq 1$ periods. (Although we assume a discrete-time setting in the analysis, our results can also be applied to continuous-time setting. Indeed, we will use a continuous-time setting for our numerical experiments in Section 7.) At the beginning of period t, the e-tailer posts the price vector $p^t = (p^t_k)$ for K products. For each location $j \in [J]$, the price vector induces a demand vector $D^t_j(p^t) = (D^t_{jk}(p^t))$ with rate vector $\lambda_j(p^t) = (\lambda_{jk}(p^t))$, where $\lambda_j(p_t) = \mathbb{E}\left[D^t_j(p^t)\right]$. (For convenience, we assume stationary rate functions. Our results can also be generalized to the case of non-stationary rates.) Demands across different periods are assumed to be independent, but can be correlated among different products within the same period. Moreover, as is common in the literature, we allow at most one customer's arrival in each period across all demand locations, i.e., $\sum_{j=1}^{J}\sum_{k=1}^{n} D^t_{jk}(p^t) \leq 1$. This is without loss of generality since we can always slice the selling season fine enough so that at most one customer arrives in each period across all locations. The quantity $\lambda_{jk}(p^t)$ can thus be interpreted as the purchase probability of product k from demand location j in period t. We will also use $\lambda^{tot}(p) = (\sum_{j=1}^{J}\lambda_{jk}(p))_{k=1}^{K}$ to denote the total purchase probability, or aggregate demand rate over all locations. Our model implicitly assumes that a customer only purchases at most one product at a time. (The case where customers purchase multiple products at the same time is challenging to analyze, even from the perspective of pure fulfillment decisions, see Jasin and Sinha 2015. We leave this for future pursuit.)

A common feature of e-commerce retail is that customers from all all demand locations observe a same price vector p^t from the same website at the same point of time. Compared to brick-and-mortar retailers where prices could be different across different physical stores, this distinct feature limits the e-tailer's degree of freedom in controlling demand intensity from multiple locations. (Technically, the e-tailer can set different prices to different customers at the same point of time according to their profiles retrieved from cookies. However, such practice may cause severe adverse effect since (1) it will lead to customer's unfair perception, psychological resistance, negative word-of-mouth, and brand switching (Zhan and Lloyd 2014), and (2) it is commonly considered as unethical if not unlawful (Reid 2014).) Indeed, this is also the very feature that makes the analysis of JPF in e-commerce setting more challenging than the typical setting in RM model. (See Section 4 for more discussions.) For each location $j \in [J]$, let $R_j^t(p^t) := (p^t)^\top D_j^t(p^t)$ denote the realized revenue in period t, where $(p^t)^\top$ indicates the transpose of p^t. We call $r_j(p^t) = \mathbb{E}\left[R_j^t(p^t)\right]$ the revenue rate for location j in period t. We use $\mathbf{J}_f$ denote the $K \times K$ Jacobian matrix for any $f = (f_1, f_2, \ldots, f_K) : \mathbb{R}^K \to \mathbb{R}^K$, i.e., $\mathbf{J}_f(x) = [(\nabla f_1(x))^\top; \ldots; (\nabla f_K(x))^\top]$ where $\nabla f_k(x)$ is the gradient of f_k at x. Let $\Omega_p := \otimes_{k=1}^K [p_l, p_u] \subset \mathbb{R}^K$ and $\Omega_\lambda \subset \mathbb{R}^K$ denote the convex and compact sets of feasible prices and demand rates, respectively. (Without loss of generality, we assume that the domain of price and demand rates at all locations are the same.) To facilitate our analysis, we make the following assumptions on the underlying demand and revenue rate functions for all $j \in [J]$:

A1. The demand rates $\lambda_j(p) : \Omega_p \to \Omega_\lambda$ and $\lambda^{tot}(p) : \Omega_p \to \otimes_{k=1}^K [0, 1]$ are invertible, twice-differentiable and monotonically decreasing in its individual argument.

A2. The revenue rates $r_j(p)$ is continuous and strictly unimodal with interior maximizer.

A3. For all $p \in \Omega_p$, the absolute eigenvalues of $\mathbf{J}_{\lambda^{tot}}(p)$ are bounded from below, whereas the absolute eigenvalues of $\nabla^2 r_j(p)$ are bounded from above.

Assumptions A1 and A2 are standard regularity conditions assumed in the RM literature (see similar assumptions in Gallego and van Ryzin 1997). The first part of A3 is a natural consequence of the invertibility of demand function; the second part of A3 is easily satisfied, especially for a compact pricing decision region. Both of them have been assumed in the dynamic pricing literature (e.g., Wang et al. 2014, Chen et al. 2015). It can be easily shown that Assumptions A1 - A3 are satisfied by a broad class of demand functions such as linear, exponential and logit demand models. Note that we do *not* assume that the revenue rate is concave when viewed as a function of demand rate instead of price, which is a critical assumption in most existing studies on dynamic pricing. As will be discussed in Section 4, we are able to sidestep the necessity of such assumption by a novel deterministic formulation of the original stochastic problem.

After a customer in location j makes a purchase of product k, the e-tailer chooses an FC i from which the order should be fulfilled immediately. We do not allow any deliberate delay in shipment for further savings in cost, since it is in itself is a complex research problem and beyond the scope of this work; see Xu et al. (2009) for further discussions on the same assumption. The shipping cost of product k from FC i to location j is $c_{ijk} \geq 0$. Let $X^t_{ijk} \in \{0,1\}$ denote the e-tailer's decision to fulfill an incoming order for product k from location j in period t using the inventory available at FC i. We assume that FC i carries $C_i = (C_{ik}) \succeq \mathbf{1}$ units of initial inventory before the selling season starts and no replenishment occurs during the selling season. (We use $\mathbf{1}$ to denote a column vector with proper dimension whose entries are all ones, and $a \succeq b$ to denote $a_i \geq b_i$ for any vectors a, b with the same dimension). The assumption on no replenishment opportunity is commonly made in previous works on dynamic fulfillment optimization (e.g., Xu et al. 2009, Acimovic and Graves 2014, and Jasin and Sinha 2015). The justifications are as follows: (1) we can interpret our selling season as the time window between two replenishments and we focus on the tactical instead of strategic decisions; and, (2) the impact of stockout can be accounted for as explained shortly. For ease of exposition, we define a *fictitious* FC 0 that has an infinite amount of initial inventory, i.e., $C_0 = +\infty \cdot \mathbf{1}$), and shipping costs set by us at $c_{0jk} := \max\{2\max_{i\in[I]} c_{ijk}, p_u\}$ for all j,k. The formulation of FC 0 serves the purpose of backup facility when certain product is depleted at all real FCs, and technically guarantees that there is always a feasible solution to our problem. In practice, the e-tailer may also decide to simply announce that the product is unavailable when it is depleted at all real FCs; in this case, the cost of shipping from FC 0 can be interpreted as the cost of lost sales. Our analysis does *not* depend on the specific cost of shipping from FC 0. For the purpose of this work, we set the cost to be no smaller than both the maximum revenue of a single product and all the other fulfillment options simply to emphasize the undesirability of fulfilling from FC 0.

In addition to having to make the pricing and fulfillment decisions, the e-tailer also needs to satisfy several constraints. First, any arriving order in period t must be fulfilled in the same period (i.e., no backorder or intentionally delayed shipment) by a unit of inventory at a certain FC. Second, the number of orders each FC fulfilled throughout the selling season cannot exceed the initial inventory level at that FC. The e-tailer's objective is to maximize the total expected profit, which is defined as total expected revenues minus total expected fulfillment costs. We can write the optimal control formulation of JPF problem as follows:

$$
\begin{aligned}
(\textbf{JPF}) \qquad \mathcal{J}^* := \max_{\{p^{t,\pi}, X^{t,\pi}\}\in\Pi} \ & \mathbb{E}^\pi\left[\sum_{t=1}^{T}\sum_{j=1}^{J}\left(p^{t,\pi}\right)^\top D_j^t(p^{t,\pi}) - \sum_{t=1}^{T}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}X_{ijk}^{t,\pi}\right] \\
\text{s.t.} \quad & \sum_{i=0}^{I} X_{ijk}^{t,\pi} = D_{jk}^t\left(p^{t,\pi}\right), \quad \forall j,k,t && (1) \\
& \sum_{t=1}^{T}\sum_{j=1}^{J} X_{ijk}^{t,\pi} \le C_{ik}, \quad \forall i,k && (2) \\
& p^{t,\pi} \in \Omega_p,\ X_{ijk}^{t,\pi} \in \{0,1\}, \quad \forall i,j,k,t && (3)
\end{aligned}
$$

where Π is the set of all non-anticipating policies and the constraints must hold almost surely. For notational brevity, we will suppress the dependency on π whenever the heuristic used is clear from the context.

Remark 1. In practice, e-tailers usually offer different options for delivery speed. Our modeling framework is general enough to cover this extra layer of complexity. Consider the original JPF model, with an addition of L different shipping options. Each shipping option ℓ of product k requires a nominal fee $s_{\ell k} \ge 0$ paid by consumers and induces a cost of $c_{ijk\ell}$ for e-tailers to adopt this option to ship a single unit of product k from FC i to location j. We can define the random demand as $D_{jk\ell}^t(p) \in \{0,1\}$ with mean $\lambda_{jk\ell}(p)$. Since the e-tailer now can collect additional nominal fee for shipping, the revenue rate is therefore $r_{jkl}(p) = (p_k + s_{\ell k}) \cdot \lambda_{jk\ell}(p)$. Lastly, we use $X_{ijk\ell}^t \in \{0,1\}$ denote the e-tailer's decision to fulfill an incoming order for product k from location j with option ℓ in period t using the inventory available at FC i. The optimal control of JPF problem with shipping options (JPF-S) can be formulated similarly as the original JPF problem. Although our exposition in the remainder of this paper is based on the original JPF formulation, all the results can be easily generalized to the case of JPF-S.

4. A Deterministic Approximation of JPF

In practice, the magnitude of demand intensity faced by an e-tailer is often high, especially during holiday seasons. (According to CNN 2015, Amazon.com sold 398 items per second during its global shopping event exclusively for Amazon Prime members on July 15, 2015.) This translates into the need for e-tailers to make fast real-time decisions, both in terms of pricing and fulfillment decisions. This requirement, together with the well-known curse of dimensionality of dynamic programming, makes solving JPF optimally practically infeasible. In the RM literature where a similar problem is encountered, many researchers turn their attention to develop heuristics that are both easy to implement and have a provably good performance under well-defined metrics. One heuristic that has drawn a lot of attention, most notably because of its practical appeal, is based on a deterministic relaxation of the original optimal control problem, where all the random variables are replaced by their means. (This approach is also called the *Certainty Equivalent* (CE) approach in the wider operations research literature, e.g., Ciocan and Farias 2012.) The benefits of such relaxation in RM literature, under some proper conditions, are threefold: (1) the resulting deterministic optimization turns out to be a concave maximization problem and is much easier to solve than the original stochastic problem; (2) its solution serves naturally as a simple heuristic; and, (3) its optimal value serves as an upper bound for the optimal control problem. Consequently, when analyzing the performance of *any* feasible policy, it suffices to benchmark it against the optimal value of the deterministic relaxation. To mimic this idea, let us first consider the following deterministic formulation of JPF, which we call *Deterministic JPF* (DJPF):

$$
\begin{aligned}
(\mathbf{DJPF}) \qquad \mathcal{J}^D := \max_{\{p^t, x^t\}} \; & \sum_{t=1}^{T}\sum_{j=1}^{J} r_j(p^t) - \sum_{t=1}^{T}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk} x_{ijk}^t \\
\text{s.t.} \; & \sum_{i=0}^{I} x_{ijk}^t = \lambda_{jk}(p^t), \quad \forall j,k,t \qquad (4) \\
& \sum_{t=1}^{T}\sum_{j=1}^{J} x_{ijk}^t \le C_{ik}, \quad \forall i,k \qquad (5) \\
& p^t \in \Omega_p, \; x_{ij}^t \in [0,1] \qquad (6)
\end{aligned}
$$

The optimal solution of DJPF has a natural interpretation: p^t can be used as the posted price vector in period t and $x_{ijk}^t/\lambda_{jk}(p^t)$ can be used as the probability of fulfilling an order of product k from location j in period t using an inventory in FC i. The important question is whether this is a good heuristic in comparison to the optimal one; if so, in what sense. It should be noted that one of the key elements in proving the near-optimality of CE-type heuristic in a typical RM literature is the fact that the optimal value of the deterministic relaxation is an upper bound of the optimal value of the original stochastic problem. Under the standard assumptions that demand rate is invertible in price and revenue is concave in demand rate (revenue does not have to be concave in price), the typical proof of this fact proceeds in two steps: First, re-write the original problem using demand rate instead of price as the decision variable; and second, apply Jensen's inequality to the objective function. When applied to JPF, unfortunately, this two-step procedure do not yield DJPF. This is so because, in any given period, the price vector observed by customers in all locations are the same, which results in new non-linear constraints that cannot be easily transformed into deterministic constraints by standard techniques. If demand rates are linear in prices, then DJPF is indeed a proper deterministic relaxation of JPF and $\mathcal{J}^* \leq \mathcal{J}^D$. In general, even if revenue rate is concave in demand rate, it is possible that $\mathcal{J}^* > \mathcal{J}^D$. This means that the performance of a heuristic derived directly from the solution of DJPF, if it is intended to mimic $\mathcal{J}^D$ at all, may in fact deviate a lot from $\mathcal{J}^*$ unless $\mathcal{J}^* \approx \mathcal{J}^D$ (see numerical results in Section 7). Motivated by the preceding discussions, in this paper, we will use an alternative deterministic formulation based on the idea of price discretization. We will show that it is possible to construct a deterministic optimization problem whose optimal value is at most $\epsilon > 0$ smaller than $\mathcal{J}^*$. We will then use this alternative deterministic formulation to construct our heuristics. (That said, our approach in this paper can also be used in combination with DJPF if the e-tailer prefers to solve DJPF instead of our proposed formulation.)

An Approximate Linear Program. We start by selecting M different price vectors $q_1, \ldots, q_M \in \Omega_p$. We will describe the precise construction of the price vectors shortly. Let $q = (q_m)_{m=1}^M$ denote the set of our discrete price vectors and $\alpha^t = (\alpha_1^t, \ldots, \alpha_M^t)$ denote a weight vector whose entries are all non-negative and sum up to one. For a fixed discretization set q, consider the following *Approximate Linear Program* (ALP):

$$(\mathbf{ALP}) \qquad \mathcal{J}^{ALP} := \max_{\{\alpha^t, x^t\}} \sum_{t=1}^T \sum_{j=1}^J \sum_{m=1}^M \alpha_m^t r_j(q_m) - \sum_{t=1}^T \sum_{i=0}^I \sum_{j=1}^J \sum_{k=1}^K c_{ijk} x_{ijk}^t$$

$$\text{s.t.} \quad \sum_{i=0}^I x_{ijk}^t = \sum_{m=1}^M \alpha_m^t \lambda_{jk}(q_m), \quad \forall j,k,t \tag{7}$$

$$\sum_{t=1}^T \sum_{j=1}^J x_{ijk}^t \leq C_{ik}, \quad \forall i,k \tag{8}$$

$$0 \leq x_{ijk}^t \leq 1, \quad \forall i,j,k,t \tag{9}$$

$$\sum_{m=1}^{M} \alpha_m^t = 1, \quad \alpha_m^t \geq 0, \forall m, t \tag{10}$$

Some comments are in order. First, the solution of ALP can be used to construct a joint pricing and fulfillment heuristic: In period t, we apply price vector q_m with probability α_m^t and fulfill an order of product k from location j using an inventory in FC i with probability $x_{ijk}^t / \sum_{m=1}^M \alpha_m^t \lambda_{jk}(q_m)$ (since FC 0 has infinite inventory, ALP always has a solution). We will formally present this heuristic and its performance in Section 5. Second, if we include the optimal prices solution from DJPF in the discretization set q, it is not difficult to see that $\mathcal{J}^D \leq \mathcal{J}^{ALP}$. Thus, one can view ALP as a generalization of DJPF that allows price vector to be sampled from a multi-point distribution instead of a singleton. Third, since ALP is an LP and demand rates are stationary, it is not difficult to see that there exists a stationary optimal solution satisfying $x_{ijk}^t = x_{ijk}^1$ and $\alpha_m^t = \alpha_m^1$ for all t. (Let (x_{ijk}^t) and (α_m^t) denote a pair of optimal solution of ALP. Define: $x_{ijk}^{t*} = \frac{1}{T}\sum_{s=1}^T x_{ijs}$ and $\alpha_m^{t*} = \frac{1}{T}\sum_{s=1}^T \alpha_m^s$. It is not difficult to check that (x_{ijk}^{t*}) and (α_m^{t*}) are also optimal for ALP.) Without loss of generality, throughout this paper we will be working with a stationary optimal solution of ALP, which is simply denoted as $\mathbf{x}^* := (x_{ijk}^*)$ and $\alpha^* := (\alpha_m^*)$. We will also assume that $\alpha_m^* > 0$ for all $m \in [M]$, since if $\alpha_m = 0$ for some m, we can simply delete q_m^* from the set q^* without affecting any of the decisions on α_m^* and x_{ijk}^*. The following lemma tells us that there exists a set of discrete price vectors q such that $\mathcal{J}^* - \mathcal{J}^{ALP} \leq \epsilon$. This means that JPF can be well-approximated by ALP, at the cost of increased computational complexity.

LEMMA 1. *Given $\epsilon > 0$, under assumptions A1, there exists a discretization q such that*

$$\mathcal{J}^* - \mathcal{J}^{ALP} \leq \epsilon$$

In proving Lemma 1, we first define a specific set of dicrete price vectors that forms a uniform grid on Ω_p and show that it indeed satisfies the approximation guarantee. Formally, given a positive integer M, we divide the feasible set $[p_l, p_u]$ into $\lfloor M^{1/K} \rfloor$ sub-intervals of equal length and let $\mathcal{Q}^* = \{p_l + (i - \frac{1}{2})(p_u - p_l)/(\lfloor M^{1/K} \rfloor + 1)\}_{i=1}^{\lfloor M^{1/K} \rfloor}$ be the set of mid-points of all these sub-intervals. We then define our *uniform grid* as the set of all possible K-permutations of $\mathcal{Q}^*$, i.e., $q^* = \{(p_1, \ldots, p_K) \in \Omega_p : p_k \in \mathcal{Q}^* \ \forall k \in [K]\}$. The number of required grid points could be large for a fine approximation (roughly proportional to ϵ^{-K} if demands are correlated among products, $K\epsilon^{-1}$ if demands are independent); however, both of our heuristics only solve it once before the selling season and is, therefore, computationally feasible. From a practical perspective, e-tailers often work with a predetermined finite set of discrete prices (see e.g. Section 5.2.1.3 in Talluri and van Ryzin 2006 and Cohen et al. 2014). In this context, our main result in this paper can be seen as a way to further exploit the given set to maximize total expected profits. Lemma 1 provides a theoretical justification that this type of approximation well approximates the original stochastic control problem (for a sufficiently fine discretization). Although our heuristics can be applied in combination with any price discretization q^*, in the remaining of this paper we will always use the set of uniform grids discussed above for consistency.

Asymptotic Regime and Performance Measure. In the sequel, we will use the optimal value of ALP as the benchmark to evaluate the theoretical performance of our heuristics. Motivated by the large volume of sales faced by e-tailers, and for the purpose of theoretical performance analysis, we will consider a sequence of JPFs and ALPs where both the length of selling season and the amount of initial inventories are scaled proportionally by a factor of θ while keeping all the other parameters unchanged. More specifically, in the θ^{th} problem, the length of selling season is given by $T(\theta) = \theta T$ and the amount of initial inventories in FC i is given by $C_i(\theta) = \theta C_i$. Since we only allow at most one new arrival in each period, increasing the selling season by θ is equivalent to multiplying the number of potential demands by θ. In other words, in the prescribed *asymptotic* setting, we essentially scale both the potential demands and initial inventories proportionally. Naturally, we shall interpret the scaling parameter θ as the *size* of the problem.

Asymptotic analysis allows us to study the theoretical performance of a given heuristic with respect to a certain benchmark (e.g., the optimal policy) without having to actually compute the exact solution of the proposed benchmark, which can be difficult. Although there is no theoretical guarantee that a heuristic that performs well in asymptotic setting (e.g., large demand and large inventory setting) will also perform well in non-asymptotic setting, existing works in the literature utilizing this approach (e.g., in RM (Gallego and van Ryzin 1994, 1997), in inventory management (Huh et al. 2009, Xin and Goldberg 2014), and in queueing (Harrison 1998, Ata and Kumar 2005)) have found that heuristics that perform well in asymptotic setting tend to also perform sufficiently well, if not extremely well, in non-asymptotic setting. This provides another motivation for asymptotic analysis.

But, most importantly, a heuristic developed in this manner can also be used as a baseline policy on which more sophisticated heuristics can be developed.

Let $\mathcal{J}^*(\theta)$ and $R^\pi(\theta)$ denote the total expected profits collected by the optimal policy and a specific heuristic $\pi \in \Pi$ on a problem with size θ, respectively; also, let $\mathcal{J}^{ALP}(\theta)$ denote the optimal value of ALP with size θ. We will use the *loss* of heuristic π for large θ as our performance measure, which is defined as $\mathcal{L}^\pi(\theta) := \mathcal{J}^{ALP}(\theta) - \mathbb{E}[R^\pi(\theta)]$.

5. First Heuristic: Randomized Pricing and Fulfillment

In this section, we describe a simple non-adaptive heuristic motivated by ALP and discuss its asymptotic performance. Let $\sigma_k^t : [J] \to [I] \cup \{0\}$ denote the fulfillment assignment for period t, i.e., $\sigma_k^t(j) = i$ indicates that we fulfill an order of product k from location j in period t from FC i. Our first heuristic uses the solution of ALP directly to construct a randomized heuristic. Note that, for a fixed set of discrete price vectors q^*, α^* and $\mathbf{x}^*$ are the optimal sampling vector and fulfillment vector given by ALP. The idea behind our first heuristic is to sample a price vector p^t from q^* according to α^*, and sample the fulfillment assignment σ^t according to $\mathbf{x}^*$. Let C_i^t denote the inventory level in FC i at the beginning of period t. We formally define our first heuristic below.

Randomized Pricing and Fulfillment Heuristic (RPF)

1. **Initialization**: Fix a discretization q^* and solve ALP to get α^*, $\mathbf{x}^*$.
2. During period $t \geq 1$, do:
   a. Sample $p^t = q_m^*$ with probability $\mathbb{P}\{p^t = q_m^*\} = \alpha_m^*$ and apply p^t.
   b. Sample $\sigma_k^t(j)$ with probability $\mathbb{P}\{\sigma_k^t(j) = i\} = y_{ijk}^* := x_{ijk}^* / \sum_{i=0}^I x_{ijk}^*$.
   c. If there exists a $(j,k) \in [J] \otimes [K]$ such that $D_{jk}^t = 1$, do:
      i. If $C_{\sigma_k^t(j),k}^t > 0$, fulfill the order from FC $\sigma_k^t(j)$ and update $C_{\sigma_k^t(j),k}^{t+1} = C_{\sigma_k^t(j),k}^t - 1$;
      ii. Otherwise, fulfill the order from FC 0.

The following theorem characterizes the performance of the RPF heuristic.

THEOREM 1. *Let q^* be the uniform price grids discussed in Section 4. There exists a constant $\Psi_1 > 0$ independent of $\theta \geq 1$ such that*

$$\mathcal{L}^{RPF}(\theta) \leq \Psi_1 \sqrt{\theta}.$$

Some comments are in order. First, it is not difficult to show that $\mathcal{J}^{ALP}(\theta)$ is an upper bound of total expected profits under *any* feasible joint pricing and fulfillment policy that restricts $p^t \in q^*$ for all t. Since $\mathcal{J}^{ALP}(\theta) = \theta \mathcal{J}^{ALP}$ is of order θ, Theorem 1 tells us that RPF is asymptotically optimal relative to $\mathcal{J}^{ALP}(\theta)$ among that class of policies (i.e., the average loss per period of RPF is of order $1/\sqrt{\theta}$, which is small for large θ). Moreover, it can be shown that the above bound is tight: For some problem instances, there exists a constant $\Psi_1' > 0$ independent of $\theta \geq 1$ such that $\mathcal{J}^{ALP}(\theta) - \mathbb{E}[R^{RPF}(\theta)] \geq \Psi_1'\sqrt{\theta}$ (see Remark 2 in Jasin 2014 for an argument for the simplistic example where $I = J = K = 1$). This means that Theorem 1 completely characterizes the asymptotic performance of RPF, in general.

Second, although RPF is asymptotically optimal, a heuristic that has a stronger performance guarantee than $\sqrt{\theta}$ is still highly desirable. For one thing, the bound in Theorem 1 is only asymptotic in nature, which means that the performance of RPF may not be too satisfactory if θ is not sufficiently large (we will test this using numerical experiments in Section 7). Another reason is the relatively thin operating margin of e-commerce retailing, as discussed in Section 1. This underscores the importance of earning (or saving) as many dollars as possible. The important question is how to construct a heuristic that maintains the tractability of RPF while at the same time significantly improves its performance guarantee. One simple idea is to re-optimize the ALP at the beginning of every period by replacing the initial inventory level C_i with the current on-hand inventory level, denoted by C_i^t. Per our brief discussions in Section 2, re-optimization-based heuristic has been shown to be effective in the broad literature of dynamic optimization. However, it is not always practically feasible. The bottleneck in our setting is the number of price vector discretizations, which can be exponential in the number of products (with a 5-point discretization for each product and a total of 10 products, we have a total of $5^{10} \approx 10^7$ price vector discretizations). Motivated by this, in this paper, we will not focus on the heuristic that fully re-optimizes ALP. Instead, in the next section, we will develop a novel re-adjust-and-re-optimize heuristic based on the idea of combining autonomous price adjustment with re-optimization of only the fulfillment part of ALP.

6. Second Heuristic: Re-adjust and Re-optimize Pricing and Fulfillment

We now discuss a modification of RPF that adaptively adjusts the discretization set q^* and updates the fulfillment vector $\mathbf{x}^*$ every period. An important feature of the proposed heuristic is that although the prices and fulfillment probabilities are still decided jointly at the beginning of the selling season via solving ALP, their adjustments during the selling season are done almost separately through a two-stage process: (1) We first adjust the discretization set q^* based on the observed demand realizations and the aggregate total of remaining inventories in all FCs at the beginning of period t; (2) we then update the fulfillment probabilities by re-optimizing the fulfillment part of ALP (we will define it shortly), which has a much smaller number of variables compared to the full ALP.

In the proposed heuristic, price is used as a lever to compensate for the fluctuation in demand realizations: To adapt with the total remaining inventories, the seller raises future prices if the realized demands in the current period are higher than expected and drops the prices otherwise. However, changing q^* directly affects the demands in *all* J locations, not just one. This calls for the second-stage adjustment on the fulfillment assignment which must be carefully calculated so as not to favor the correction in only a few locations. More specifically, we solve another LP to update our fulfillment assignment accordingly. Our main result in this section is to show that, under some conditions, a scheme that partially decouples the adjustment in prices from the adjustment in fulfillment and guarantees a significant improvement over RPF exists.

We start with defining some notations that will be useful for our exposition. Let $C^t := (C_i^t)$ denote the vector of inventory level at the beginning of period t. Given the new discretization set $q^t = (q_m^t)$ (by definition, we have $q^1 = q^*$), we update the fulfillment probabilities by solving the following *Fulfillment LP* (FLP):

$$\mathbf{FLP}^t(q^t, C^t) := \left\{ \min_{x_{ijk} \geq 0} \mathbf{c}^\top \mathbf{x}, \; : \; \sum_{i=0}^{I} x_{ijk} = \sum_{m=1}^{M} \alpha_m^* \lambda_{jk}\left(q_m^t\right), \; \sum_{j=1}^{J} x_{ijk} \leq \frac{C_{ik}^t}{T-t+1} \right\}, \quad (11)$$

where $\mathbf{c}$ and $\mathbf{x}$ denote the cost and fulfillment vectors, respectively. For notational brevity, we will often write $\mathrm{FLP}(q^t, C^t)$ simply as FLP^t whenever the values of q^t and C^t used are clear from the context. Note that we drop the constraints $x_{ijk} \leq 1$; this is without loss of generality since $x_{ijk} \leq \sum_{m=1}^{M} \alpha_m^* \lambda_{jk}(q_m^t) \leq \sum_{m=1}^{M} \alpha_m^* \lambda_{jk}(p_l \ldots, p_l) \leq 1$. Let $y_{ijk}^t := x_{ijk}^t / \sum_{i=0}^{I} x_{ijk}^t$ and let Y_{ijk}^t be an indicator variable with $Y_{ijk}^t = 1$ if $\sigma_k^t(j) = i$ and 0 otherwise. (By definition, $\mathbb{P}(Y_{ijk}^t = 1) = y_{ijk}^t$.) In other words, y_{ijk}^t is the conditional probability of using FC i to fulfill an order of product k from location j conditioning on such order being observed. Define:

$$\Delta C_{ik}^t := \sum_{j=1}^{J} \left[X_{ijk}^t - y_{ijk}^t \left(\sum_{m=1}^{M} \alpha_m^* \lambda_{jk}\left(q_m^t\right) \right) \right]. \quad (12)$$

Note that ΔC_{ik}^t can be interpreted as the size of aggregate randomness in inventory consumption of product k at FC i. To be precise, ΔC_{ik}^t should have been written as a function of p^t and q^t. We suppress this notational dependencies for the sake of brevity. Let $\Delta C_i^t = [\Delta C_{i1}^t; \Delta C_{i2}^t; \ldots; \Delta C_{iK}^t]^\top$. We are now ready to define our second heuristic.

Re-adjust and Re-optimize Pricing and Fulfillment Heuristic (R^{2}PF)

---

1. **Initialization**: Fix discretization q^* and solve ALP to get α^*, $\mathbf{x}^*$.
   Define $q^1 = q^*$ and $\hat{\mathbf{x}}^1 = \mathbf{x}^*$.
2. During period $t \geq 1$, do:
   a. **Adjust Price:** For each m, define q_m^t such that
   $$\lambda^{tot}\left(q_m^t\right) := \underset{\otimes_{k=1}^K [0,1]}{\text{proj}} \left[\lambda^{tot}(q_m^*) - \frac{1}{M\alpha_m}\left(\sum_{i=0}^{I}\sum_{s=1}^{t-1}\frac{\Delta C_i^s}{T-s}\right)\right].$$
   b. **Adjust Fulfillment:** Set $\hat{\mathbf{x}}^{t+1}$ equals to the optimal solution of $\mathbf{FLP}^t(q^{t+1}; C^t)$.
   c. Sample p^t with probability $\mathbb{P}\{p^t = q_m^t\} = \alpha_m^*$ and apply p^t.
   d. Sample $\sigma_k^t(j)$ with probability $\mathbb{P}\{\sigma_k^t(j) = i\} = y_{ijk}^t := \hat{x}_{ijk}^t / \sum_{i=0}^I \hat{x}_{ijk}^t$.
   e. If there exists a $(j,k) \in [J] \otimes [K]$ such that $D_{jk}^t = 1$, do:
      i. If $C_{\sigma_k^t(j),k}^t > 0$, fulfill the order from FC $\sigma_k^t(j)$ and update $C_{\sigma_k^t(j),k}^{t+1} = C_{\sigma_k^t(j),k}^t - 1$;
      ii. Otherwise, fulfill the order from FC 0.

Note that, by definition, ΔC_{ik}^s is the error (or deviation) from the expected consumption of product k in FC i at period s. So, the term $\Delta C_{ik}^s/(T-s)$ can be interpreted as the portion of this error to be corrected in period $t > s$. (The term $T-s$ in the denominator indicates that the correction for the error incurred in period s is to be distributed uniformly throughout the remaining periods. Although this may not be the optimal correction mechanism, Jasin 2014 has shown in the context of dynamic pricing that it is sufficient to guarantee a very strong performance bound.) Thus, at period t, the cumulative errors (across all FCs) for product k that needs to be corrected is given by $\sum_{i=0}^I \sum_{s=1}^t \Delta C_{ik}^s/(T-s)$. Our idea is to correct these errors by perturbing the original set of discrete price vectors q^* to q^t such that the following system of balance equations holds:

$$\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^*\lambda_{jk}\left(q_m^t\right) = \sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^*\lambda_{jk}\left(q_m^*\right) - \sum_{i=0}^{I}\sum_{s=1}^{t-1}\frac{\Delta C_{ik}^s}{T-s}, \quad \forall k. \tag{13}$$

If $\lambda^{tot}(q_m^*) - (\sum_{i=0}^I \sum_{s=1}^{t-1} \Delta C_i^s/(T-s))/(M\alpha_m) \in \otimes_{k=1}^K [0,1]$, it can be easily verified that a solution to the system of non-linear equations in Step 2a is also a solution to (13). Moreover, by the invertibility of $\lambda^{tot}(\cdot)$ (Assumption A1), the system in Step 2a always has a unique solution of q^t. Although we need to perturb potentially all price vectors in q^*, the computation in Step 2a can be done in parallel. This decomposability is crucial for the time-efficiency of R^{2}PF. We want to emphasize: *Although Step 2a helps balance future demands with remaining inventories, it only does so at an* aggregate *(across all the FCs) level.* To address the potential inventory imbalance that exists across different FCs due to demand randomness and our fulfillment heuristic, another layer of adjustment is needed. We do this by re-optimizing FLP in Step 2b, parameterized with adjusted expected demands under the new discretization set q^t. This extra step is crucial for making sure that we are also minimizing total shipping costs while maximizing total revenues.

Before we evaluate the asymptotic performance of R^{2}PF, we need to first introduce a concept that will be useful for the analysis. Consider the initial transportation problem faced by the e-tailer, i.e., FLP^1. Since we assume that each customer only requests at most one product, FLP^1 can be decomposed into K transportation LPs defined as follows:

$$\mathrm{FLP}_k^1(q^*, C_k) := \left\{ \min_{x_{ijk} \geq 0} \sum_{i=0}^{I} \sum_{j=1}^{J} c_{ijk} x_{ijk} : \sum_{i=0}^{I} x_{ijk} = \sum_{m=1}^{M} \alpha_m^* \lambda_{jk}(q_m^*), \sum_{j=1}^{J} x_{ijk} \leq \frac{C_{ik}}{T} \right\}.$$

We assume without loss of generality that $\sum_{j=1}^{J} x_{ijk}^* = C_{ik}$ (otherwise, we can always define $\tilde{C}_{ik} := \sum_{j=1}^{J} x_{ijk}^*$ and replace the original initial inventory C_{ik} with $\tilde{C}_{ik}$ without changing anything else). The inventory constraints in FLP_k^1 are, therefore, all binding. Now, from the literature of transportation LP (e.g., Dantzig and Thapa 2006), we know that there is exactly one redundant constraint in every FLP_k^1. Moreover, if we delete an arbitrary constraint, the remaining constraints are always linearly independent. Let $\overline{\mathrm{FLP}}_k^1$ be the LP where we delete the inventory constraint regarding FC 0; by Theorem 2.5 in Bertsimas and Tsitsiklis (1997), $\overline{\mathrm{FLP}}_k^1$ is equivalent to FLP_k^1. We call a basic solution to FLP^1 as *DR-degenerate* ("DR" is short for de-redundancy) if and only if the corresponding basic solution to $\overline{\mathrm{FLP}}_k^1$ is degenerate for some $k \in [K]$.

The following theorem characterizes the performance of R^{2}PF.

THEOREM 2. *Let q^* be the uniform price grids discussed in Section 3. Suppose that FLP$^1(q^*, C)$ has a unique non-DR-degenerate optimal solution. There exists a constant $\Psi_2 > 0$ independent of $\theta \geq 1$ such that*

$$\mathcal{L}^{R^2PF}(\theta) \;\leq\; \Psi_2(1 + \log\theta).$$

Some comments are in order. First, since $\mathrm{R}^2\mathrm{PF}$ may use different discretization sets in different periods, $\mathcal{J}^{ALP}(\theta)$ is not necessarily an upper bound for $\mathbb{E}[R^{R^2PF}(\theta)]$. However, given that the expected loss of RPF relative to $\mathcal{J}^{ALP}(\theta)$ is of order $\sqrt{\theta}$, the bound in Theorem 2 is still useful because it shows that $\mathrm{R}^2\mathrm{PF}$ guarantees a significant improvement over RPF, at least asymptotically. Second, the fact that $\mathrm{R}^2\mathrm{PF}$ significantly improves RPF is quite surprising. For one thing, aside from the definition $\{\Delta C_{ik}^t\}$, the adjustment formula in Step 2a is pretty much independent of the actual inventory distribution across I FCs at the beginning of period $t+1$ — they only depend on the aggregate errors. Moreover, although it is known in the literature that frequent re-optimizations has a potential to significantly improve performance (see Section 2), it matters what is being re-optimized. In the case of $\mathrm{R}^2\mathrm{PF}$, the fulfillment LP takes as its input the new discretization set that is adjusted almost independently of the current inventory distribution and how it would affect total shipping costs. It is, thus, not immediately clear that frequent re-optimizations of the fulfillment LP updated in this manner still yields the level of improvement that we want. Fortunately, we show that the proposed combination of aggregate demand adjustments and fulfillment re-optimizations still give a significant improvement over RPF. Lastly, we want to emphasize that the non-DR-degeneracy assumption only applies to the initial FLP^1 and is not required for the subsequent FLP^t for all $t \geq 2$. Similar conditions have been used in other works that study the performance of re-optimization-based heuristic with deterministic relaxation being an LP, e.g., Jasin and Kumar (2012, 2013), and Johnson et al. (2015). Although this assumption is critical for the tractability of the proof, our numerical results in Section 7 show that $\mathrm{R}^2\mathrm{PF}$ still performs well when FLP^1 is degenerate.

7. Numerical Experiments

We now conduct two numerical experiments to illustrate the performance of the proposed heuristics in comparison to some natural benchmarks. The setting of our numerical study is placed in the continental United States. We set $I=6$ and $J=15$ (i.e., the e-tailer has six FCs serving fifteen different demand locations) and select our fifteen demand locations to be the fifteen largest metropolitan statistical areas (MSAs) estimated by U.S. Census Bureau (2014a). The demand process is generated as follows: We first generate a sequence of Poisson arrivals with arrival rate from location j to be $\gamma_j = \textit{pois-rate} \times \textit{mkt-share}_j$. Specifically, $\textit{pois-rate} \in (0,1]$ denote the probability of a new arrival and $\textit{mkt-share}_j$ is the conditional probability that this arrival comes from region j. We set the default value of *pois-rate* to be 0.9 and $mkt\text{-}share_j$ to be the ratio of total population in the j^{th} largest MSA and $\sum_j mkt\text{-}share_j$. A customer arriving from location j makes a purchase with probability $\exp(A_j + B_j p)$. The demand parameters are chosen as follows: We first set "baseline" demand parameters A_1 and B_1. For all $j \geq 2$, we set $A_j = \frac{income_1}{income_j} \times A_1$ and $B_j = \frac{income_1}{income_j} \times B_1$, where $income_j$ represents the medium household income of the j^{th} largest MSA, as reported in U.S. Census Bureau (2014b). Since we want $\exp(A_j + B_j p) \leq 1$ for all $p \in \Omega_p$, we set A_j's to be vectors with negative components, and B_j's to be diagonally dominated matrices with negative diagonal components. (The baseline parameters are generated to satisfy these constraints.) By our construction of A_j and B_j, a customer from a demand region with higher income is more likely to make a purchase compared to a customer from a demand region with lower income. For simplicity, we normalize T to 1. This means that the scaling factor θ is the length of selling season and can be immediately interpreted as the size of potential market. For both experiments, we use the value of θ ranging from 100 to 2,000 where, in our setting, $\theta = 100$ corresponds to the case where we have about 2 units of inventory for each product in each FC and $\theta = 2,000$ corresponds to the case with about 40 units of inventory for each product in each warehouse (see the discussions how we set the value of C_{ik} below). We intentionally choose these numbers to highlight the performance of our heuristics in non-asymptotic setting. (In reality, an e-tailer can easily sell this amount of inventory within a few hours during peak season.)

The logistic networks of both experiments consist of six FCs selected from the list of the most efficient warehouses (in terms of possible transit lead-times) in the U.S., as reported by Chicago Consulting (2013). The outbound shipping costs of a single-item package from different FCs to different locations are calculated using the cost equation estimated in Section EC.3 in Jasin and Sinha (2015), assuming that each package weighs exactly one pounds. As for the fictious stores, per Section 3, their costs are calculated as $c_{0jk} := \max\{2 \max_{i \in [I]} c_{ijk}, p_u\}$ for all j, k. The average shipping cost over all FC-MSA pair is \$9.55. We, therefore, set the feasible price range to be \$100 and \$250, since the annual outbound transportation costs as a percentage of net sales typically varies between 4% to 10% (Tompkins Supply Chain Consortium 2012). (In reality, the ratio between price range and shipping cost highly varies with the type of product; our choice above at least guarantees that the relative magnitude between revenue and cost is practical.) The initial inventory level is set to be slightly higher than the number of total arrivals, which reflects the common reality where firm stocks neither too low such that the induced demand has to be really scarce, nor too high as if there is no inventory constraint at all. Specifically, we first match between FCs and MSAs such that (1) each FC serves five MSA, (2) each MSA is served by 2 FCs, and (3) the total mileage between all the assigned FC-MSA pairs is minimized. For each MSA j, we define $\hat{\lambda}_j = \textit{pois-rate} \times \textit{mkt-share}_j \times 0.9$ and let each of the two FCs serving MSA_j to fulfill a portion of the $\hat{\lambda}_j$, where the portion is decided by a random number drawn uniformly from $[0.4, 0.6]$. (Our results are robust with respect to slight perturbation in the numbers 0.9, 0.4, and 0.6.) The initial inventory at each of the FC is then calculated as the sum of all the demand portions from the five MSAs it serves. The initial inventory level at the fictitious store is set to be 20 for all products, so that the available inventory will never be depleted. For a specific θ, we always round down θC_{ik}.

We denote by *RPF-m* and *R^2PF-m* the RPF and R^2PF heuristics proposed in Sections 5 and 6 respectively, where m denotes the number of discretizations we select for the price of a single product, i.e., $M = m^K$. (It is noteworthy that, in both experiments, the initial FLPs are DR-degenerate.) As a benchmark, we implement a *Full-reopt-m* heuristic, which re-optimizes an updated ALP in each period for the optimal α^t and x^t. Compared to the original ALP, the updated ALP replaces the inventory parameters C_{ik} with the current inventory level while keeping everything else the same. As discussed in Section 5, although *Full-reopt-m* potentially has a very good performance, it may not be practically feasible due to heavy computation burden. Finally, to test whether the joint pricing and fulfillment optimization indeed generates higher profit than separate pricing and optimization, we also implement a *Sep-reopt* heuristic, where the pricing decisions and fulfillment decisions are optimized separately in every period as follows: At the beginning of period t, we first compute the new price vector by solving a constrained revenue maximization problem defined as $\{\max_{p\in\Omega} \sum_{j=1}^J r_j(p) : \lambda_{jk}(p) \le \sum_{i=1}^I C_{ik}^{t-1}/(T-t),\ \forall k\}$, where the inventory of product k is aggregated over all FCs; and, then we optimize the fulfillment assignment decisions by solving a transportation LP taking the new price vector and the current inventory levels as input, defined as $\{\min_{0\le x_{ijk}\le 1}\ \sum_{i=0}^I \sum_{j=1}^J c_{ijk}x_{ijk} : \sum_{i=0}^I x_{ijk} = \lambda_{jk}(p^t)\ \forall j,k,\ \sum_{j=1}^J x_{ijk} \le C_{ik}^{t-1}/(T-t)\ \forall i,k\}$. For a specific choice of θ, we simulate all the heuristics for 100 runs to approximate their total expected profit. For each run, we first generate the arrival sequence, then use it as input for all the heuristics; this allows us to minimize the impact of randomness in demand generation.

Experiment 1

The purpose of the first experiment is to illustrate that: (1) both *RPF-m* and R^2*PF-m* perform well, even for a relatively sparse set of price discretization and without the non-DR-degeneracy assumption; and (2) even the static heuristic *RPF-m* dominates the performance of *Sep-reopt*. To do this, we choose $K = 9$ and $m = 3$, which means that $M = 3^9 = 19,683$. For each θ, we run three algorithms, namely *RPF-3*, R^2*PF-3* and *Sep-reopt*), and compares their average performances. (We do not run *Full-reopt* in this experiment because the required computation time is simply too long; see Table 3.) The results can be found in Table 1. In addition to reporting the loss (as defined in Section 4) and percentage loss, we also report the percentage improvement in total profits for both *RPF-3* and R^2*PF-3* relative to the profit of *Sep-reopt*; this helps us better understand the benefit of joint pricing and fulfillment optimization.

θ	RPFC-3			R^{2}PFC-3			Sep-reopt	
	Loss	% Loss	%Improve	Loss	% Loss	%Improve	Loss	% Loss
100	3955.73	58.22%	0.79%	3912.13	57.57%	2.33%	3977.89	58.54%
200	5614.07	41.31%	3.67%	5204.59	38.30%	8.99%	5896.56	43.39%
400	5865.35	21.58%	3.29%	4560.62	16.78%	9.61%	6543.76	24.08%
600	6695.66	16.42%	2.25%	5523.25	13.55%	5.77%	7446.25	18.26%
800	7713.76	14.19%	2.43%	6850.85	12.60%	4.33%	8821.08	16.23%
1000	9559.36	14.07%	1.71%	8376.92	12.33%	3.77%	10539.29	15.51%
1500	9729.57	9.55%	1.74%	8072.33	7.92%	3.57%	11303.72	11.09%
2000	11959.26	8.80%	1.93%	7825.35	5.76%	5.33%	14304.14	10.53%

Table 1 Performances of different heuristics with varying θ

We now make several observations. First, it is obvious that the percentage loss of *RPF-3* and R^2*PF-3* both converges to zero as θ grows large; moreover, R^2*PF-3* converges significantly faster than *RPF-3*. This empirically validates our theoretical results in Theorem 1 and 2. Second, for all $\theta \geq 200$, the performance of *RPF-3* dominates that of *Sep-reopt* by at least 1.7%. The real-time adjustment in R^2*PF-3* brings an additional 3.5% in profit. This illustrates the benefit of joint pricing and fulfillment optimization, even if the e-tailer only performs it once before the selling season with relatively sparse price discretization. Third, under our choice of parameters, $\mathcal{J}^{ALP} = 67.9494 > \mathcal{J}^{D} = 63.8075$. This illustrates our point in Section 4 that the CE-type deterministic relaxation may not be a proper benchmark for evaluating heuristics.

Experiment 2

The purpose of the second experiment is to (1) study how the number of price discretizations affects the performance of *RPF-m* and R^2*PF-m*, (2) compare the performance of R^2*PF-m* with that of *Full-reopt-m*, and (3) test the robustness of the proposed heuristic. For this experiment, we choose $K = 5$ and $m \in \{2, 5, 8\}$, which means that $M \in \{32, 3125, 32768\}$. We run *RPF-m*, R^2*PF-m*, *Sep-reopt*, and *Full-reopt-m* for different m, with an exception of *Full-reopt-8* which we do not run because its computation time is too long (see Table 3). Note that, with different choices of m, the exact value of $\mathcal{J}^{ALP}$ is different. Hence, the absolute losses with respect to $\mathcal{J}^{ALP}$ are not directly comparable among the heuristics with different number of discretizations (the percentage loss can still be used to test the robustness of the heuristics). Due to the same reason, we also report the expected profits of *Sep-reopt* instead of its losses.

Several findings can be drawn from the results of the second experiments. First, the performance of all heuristics tested in experiment 1 are robust, as the results in this experiment reclaim the findings in the first experiment. Second, for the same choice of m, R^2*PF-m* always has comparable performance with *Full-reopt-m*; moreover, as shown in Table 3, this is achieved with a significant reduction in computation time. Third, for all three heuristics that manage pricing and fulfillment decisions jointly, in general, finer discretization leads to higher profit as long as θ is large enough (compare the numbers in *%Improve* columns). Fourth, compared to the first experiment, the improvement of joint pricing and fulfillment optimization over separate pricing and fulfillment is significantly larger. This can be explained by the relatively larger gap between $\mathcal{J}^{ALP}$ and $\mathcal{J}^D$, since the expected total profit of *Sep-reopt* can at best be closest to $\mathcal{J}^D$. (Specifically, $\mathcal{J}^{ALP} = 66.1347, 69.5241, 70.5807$ for $m = 2, 5, 8$, respectively, whereas $\mathcal{J}^D = 53.3431$.) Finally, for some instances (e.g., R^2*PFC-2* when $\theta = 1500$) the loss is actually negative, which suggest that the total expected profits of R^2*PFC-2* can be greater than $\mathcal{J}^{ALP}(\theta)$. This, together with the finding in the first experiment, confirms our conjecture in Section 4 that, in general, neither $\mathcal{J}^{ALP}$ nor $\mathcal{J}^D$ is an upper bound for different heuristics.

Lastly, we report the running time of a single simulation for *Full-reopt-m* and *Full-reopt-m* with different values of m. For the same problem instance, R^2*PFC-m* runs much

θ	RPFC-2			RPFC-5			RPFC-8		
	Loss	%Loss	%Improve	Loss	%Loss	%Improve	Loss	%Loss	%Improve
100	3147.60	47.59%	110.63%	3876.79	58.62%	74.84%	3953.14	59.77%	72.55%
200	3754.57	28.39%	58.32%	4649.50	35.15%	50.70%	5335.73	40.34%	40.76%
400	3972.55	15.02%	49.07%	6833.73	25.83%	36.76%	6436.06	24.33%	41.66%
600	4587.75	11.56%	44.87%	7214.78	18.18%	40.89%	7679.50	19.35%	40.98%
800	5862.10	11.08%	46.39%	8449.77	15.97%	45.43%	7926.62	14.98%	49.38%
1000	6703.25	10.14%	47.32%	8831.96	13.35%	49.33%	9579.95	14.49%	49.62%
1500	8737.27	8.81%	45.97%	10625.81	10.71%	50.24%	11352.69	11.44%	51.27%
2000	8601.42	6.50%	48.88%	11026.35	8.34%	53.44%	13196.07	9.98%	52.99%
θ	R^2PFC-2			R^2PFC-5			R^2PFC-8		
	Loss	%Loss	%Improve	Loss	%Loss	%Improve	Loss	%Loss	%Improve
100	2538.69	36.52%	155.16%	3256.51	46.84%	124.61%	3254.96	46.82%	128.12%
200	1826.41	13.14%	92.03%	3707.53	26.66%	70.43%	4407.40	31.70%	61.15%
400	2201.77	7.92%	61.52%	2876.74	10.34%	65.32%	1954.38	7.03%	74.04%
600	3155.06	7.56%	51.42%	2865.25	6.87%	60.37%	2619.60	6.28%	63.84%
800	1262.29	2.27%	60.90%	4660.29	8.38%	58.57%	1742.11	3.13%	70.20%
1000	3039.23	4.37%	56.77%	2544.18	3.66%	66.04%	3407.57	4.90%	66.39%
1500	-761.32	-0.73%	61.23%	1674.99	1.61%	65.56%	4980.87	4.78%	62.66%
2000	991.53	0.71%	58.10%	4078.56	2.93%	62.49%	2199.75	1.58%	67.25%
θ	Full-reopt-2			Full-reopt-5			Sep-reopt		
	Loss	%Loss	%Improve	Loss	%Loss	%Improve	Profit		
100	2403.49	34.05%	165.06%	3253.78	46.10%	127.74%	1645.45		
200	1454.08	10.30%	98.30%	3728.16	26.41%	71.02%	5983.23		
400	1714.29	6.07%	64.76%	2390.11	8.47%	68.79%	15081.29		
600	2741.67	6.47%	53.20%	1693.57	4.00%	65.31%	24224.53		
800	1094.80	1.94%	61.44%	3035.18	5.38%	63.77%	32136.91		
1000	2289.27	3.24%	58.62%	2617.94	3.71%	65.95%	40340.76		
1500	-2511.69	-2.37%	63.86%	889.25	0.84%	66.85%	61977.02		
2000	-565.48	-0.40%	59.87%	2360.87	1.67%	64.60%	83064.69		

Table 2 Expected loss of different heuristic with varying θ

faster than *Full-reopt-m*. All experiments were implemented on a desktop computer with 3.40GHz Intel Core i7-3770 CPU and 8 GB of RAM.

R^2PFC-5 (K=5)	R^2PFC-8 (K=5)	R^2PFC-9 (K=3)
23.12	23.69	25.49
Full-reopt-5 (K=5)	Full-reopt-8 (K=5)	Full-reopt-9 (K=3)
992.87	10814.08	7170.38

Table 3 Typical running time (in seconds) for a single simulation for selected heuristics ($\theta = 2000$)

8. Summary

In this paper, we consider a joint pricing and fulfillment problem faced by e-commerce retailers. In this environment, the inability of the e-tailer to charge different prices for customers coming from different regions introduces a subtlety that does not permit a trivial generalization of methodology in existing literature. To cope with this, we propose a novel LP-based approximation scheme and show that it approximates the original optimal control formulation well. Motivated by the structure of the optimal solution to the approximation LP, we further propose two different heuristics that have a strong performance guarantee. In our major heuristic, we frequently adjust the pricing decision according to an autonomous error-correcting scheme and re-solving the fulfillment decision by an updated transportation LP. We believe that this heuristic has desirable features since (1) it decouples the real-time decision in pricing and fulfillment assignment, which in practice may be managed by different functions, and (2) it yields a high total expected profits without sacrificing the computation time. Our numerical results show that both heuristics not only have good theoretical performance in asymptotic setting but also work well in non-asymptotic setting with only a few units of inventory in each FC. We further believe that our analytical framework can be used to address other stochastic optimization problems abound in the broader OM context where different related decisions have to be made jointly in real-time.

References

Acimovic, J, SC Graves. 2014. Making better fulfillment decisions on the fly in an online retail environment. *M&SOM* **17**(1) 34–51.

Agatz, NAH, M Fleischmann, JA van Nunen. 2008. E-fulfillment and multi-channel distribution–a review. *Eur. J. Oper. Res.* **187**(2) 339–356.

Ahn, HS, S Jasin, P Kaminsky, Y Wang. 2015. Certainty equivalent planning for multi-product batch differentiation: Analysis and bounds. *Available at SSRN 2704595* .

Amazon.com. 2015. Amazon annual report. `http://phx.corporate-ir.net/phoenix.zhtml?c=97664&p=irol-newsArticle&ID=2133281`.

Ata, B, S Kumar. 2005. Heavy traffic analysis of open processing networks with complete resource pooling: Asymptotic optimality of discrete review policies. *Ann. App. Probab.* **15**(1A) 331–391.

Atar, R, MI Reiman. 2012. Asymptotically optimal dynamic pricing for network revenue management. *Stochastic Systems* **2**(2) 232–276.

Aydin, G, S Ziya. 2008. Pricing promotional products under upselling. *M&SOM* **10**(3) 360–376.

Becerril-Arreola, R, M Leng, M Parlar. 2013. Online retailers promotional pricing, free-shipping threshold, and inventory decisions: A simulation-based analysis. *Eur. J. Oper. Res.* **230**(2) 272–283.

Ben-Tal, A, A Goryashko, E Guslitzer, A Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Math. Program.* **99**(2) 351–376.

Bertsimas, D, DA Iancu, PA Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Math. of Oper. Res.* **35**(2) 363–394.

Bertsimas, D, JN Tsitsiklis. 1997. *Introduction to linear optimization*, vol. 6. Athena Scientific Belmont, MA.

Boyd, EA, IC Bilegan. 2003. Revenue management and e-commerce. *Manage. Sci.* **49**(10) 1363–1386.

Calafiore, GC. 2009. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.* **48**(4) 2254–2274.

Camelcamelcamel.com. 2016. Amazon price history for lodge l8sgp3 pre-seasoned cast-iron square grill pan. `http://camelcamelcamel.com/Lodge-L8SGP3-Pre-Seasoned-Cast-Iron-10-5-inch/product/B0000CF66W?context=browse`. Accessed: 2016-03-16.

Chen, J. 2014. Pricing strategies in the electronic marketplace. *Handbook of Strategic e-Business Management*. Springer, 489–522.

Chen, Q, S Jasin, I Duenyas. 2015. Real-time dynamic pricing with minimal and flexible price adjustment. *Manage. Sci.* Forthcoming.

Chicago Consulting. 2013. The 10 best warehouse network 2013. `http://www.chicago-consulting.com/10best.shtml`.

Ciocan, DF, V Farias. 2012. Model predictive control for dynamic resource allocation. *Math. of Oper. Res.* **37**(3) 501–525.

CNN. 2015. Amazon 'prime day' shattered global sales records. `http://money.cnn.com/2015/07/15/news/amazon-walmart-sales/`.

Cohen, MC, NHZ Leung, K Panchamgam, G Perakis, A Smith. 2014. The impact of linear optimization on promotion planning. *Available at SSRN 2382251* .

Dantzig, GB, MN Thapa. 2006. *Linear programming 1: Introduction*. Springer.

Dinlersoz, EM, H Li. 2006. The shipping strategies of internet retailers: Evidence from internet book retailing. *Quant. Marketing Econom.* **4**(4) 407–438.

Doğru, MK, MI Reiman, Q Wang. 2010. A stochastic programming based inventory policy for assemble-to-order systems with application to the W model. *Oper. Res.* **58**(4-part-1) 849–864.

Ferreira, KJ, BHA Lee, D Simchi-Levi. 2015. Analytics for an online retailer: Demand forecasting and price optimization. *M&SOM* **18**(1) 69–88.

Fisher, M, S Gallino, J Li. 2015. Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Available at SSRN 2547793* .

Gallego, G, GJ van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Manage. Sci.* **40**(8) 999–1020.

Gallego, G, GJ van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* **45**(1) 24–41.

Gümüş, M, S Li, W Oh, S Ray. 2013. Shipping fees or shipping free? A tale of two price partitioning strategies in online retailing. *Prod. Oper. Manag.* **22**(4) 758–776.

Harrison, JM. 1998. Heavy traffic analysis of a system with parallel servers: Asymptotic optimality of discrete-review policies. *Ann. App. Probab.* 822–848.

Huh, WT, G Janakiraman, JA Muckstadt, P Rusmevichientong. 2009. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Manage. Sci.* **55**(3) 404–420.

Jasin, S. 2014. Reoptimization and self-adjusting price control for network revenue management. *Oper. Res.* **62**(5) 1168–1178.

Jasin, S, S Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Math. of Oper. Res.* **37**(2) 313–345.

Jasin, S, S Kumar. 2013. Analysis of deterministic lp-based booking limit and bid price controls for revenue management. *Oper. Res.* **61**(6) 1312–1320.

Jasin, S, A Sinha. 2015. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Oper. Res.* **63**(6) 1336–1351.

Johnson, K, D Simchi-Levi, H Wang. 2015. Online network revenue management using thompson sampling. *Available at SSRN 2588730* .

Leng, M, R Becerril-Arreola. 2010. Joint pricing and contingent free-shipping decisions in B2C transactions. *Prod. Oper. Manag.* **19**(4) 390–405.

Maglaras, C, J Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *M&SOM* **8**(2) 136–148.

Moallemi, Ciamac C, Mehmet Saglam. 2012. Dynamic portfolio choice with linear rebalancing rules. *Available at SSRN 2011605* .

Netessine, S, S Savin, W Xiao. 2006. Revenue management through dynamic cross selling in e-commerce retailing. *Oper. Res.* **54**(5) 893–913.

Özer, Ö, R Phillips. 2012. *The Oxford handbook of pricing management*. Oxford University Press.

Plambeck, EL, AR Ward. 2006. Optimal control of a high-volume assemble-to-order system. *Math. of Oper. Res.* **31**(3) 453–477.

Reid, B. 2014. Avoiding everyday antitrust law violations. `http://www.huffingtonpost.com/brad-reid/avoiding-everyday-antitru_b_7999742.html`.

Reiman, MI, Q Wang. 2008. An asymptotically optimal policy for a quantity-based network revenue management problem. *Math. of Oper. Res.* **33**(2) 257–282.

Rigby, D. 2014. Online shopping isnt as profitable as you think. `https://hbr.org/2014/08/online-shopping-isnt-as-profitable-as-you-think`.

Secomandi, N. 2008. An analysis of the control-algorithm re-solving issue in inventory and revenue management. *M&SOM* **10**(3) 468–483.

Secomandi, N, F Margot. 2009. Reoptimization approaches for the vehicle-routing problem with stochastic demands. *Oper. Res.* **57**(1) 214–230.

Shpanya, A. 2014. Why dynamic pricing is a must for ecommerce retailers. `https://econsultancy.com/blog/65327-why-dynamic-pricing-is-a-must-for-ecommerce-retailers/`.

Sides, R, HK Hogan. 2015. Deloitte's 2015 holiday survey: Embracing retail disruption. `http://dupress.com/articles/holiday-retail-sales-consumer-survey/`.

Simchi-Levi, D, SD Wu, ZJM Shen. 2004. *Handbook of quantitative supply chain analysis: Modeling in the e-business era*, vol. 74. Springer.

Talluri, KT, GJ van Ryzin. 2006. *The theory and practice of revenue management*, vol. 68. Springer.

Tompkins Supply Chain Consortium. 2012. Supply chain metrics: Data for you to compare against. `http://www.supplychainconsortium.com/Report/GetReport.asp?ID=59`.

UPS. 2014. Pulse of the online shopper. `https://www.ups.com/media/en/2014-UPS-Pulse-of-the-Online-Shopper.pdf`.

UPS. 2016. 2016 UPS rate and service guide. `https://www.ups.com/media/en/daily_rates.pdf`.

U.S. Census Bureau. 2014a. Annual estimates of the resident population: April 1, 2010 to july 1, 2014 - United States - metropolitan and micropolitan statistical area. `http://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?src=bkmk`.

U.S. Census Bureau. 2014b. Income and poverty in the United States: 2014. `https://www.census.gov/hhes/www/income/data/`.

U.S. Census Bureau. 2016. Latest quarterly e-commerce report. `http://www.census.gov/retail/mrts/www/data/pdf/ec_current.pdf`.

Wang, Z, S Deng, Y Ye. 2014. Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Oper. Res.* **62**(2) 318–331.

Williams, D. 1991. *Probability with martingales*. Cambridge university press.

Xin, L, DA Goldberg. 2014. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. *arXiv preprint arXiv:1409.1499* .

Xu, PJ, R Allgor, SC Graves. 2009. Benefits of reevaluating real-time order fulfillment decisions. *M&SOM* **11**(2) 340–355.

Zhan, L, AE Lloyd. 2014. Customers asymmetrical responses to variable pricing. *J. Revenue Pricing Management* **13**(3) 183–198.

Electronic Companion

EC.1. Proof of Lemma 1

In what follows, we will only show the existence of a proper set q^* under the single-product setting; the argument can be easily extended to the multiple-product setting. Let $F^t : \Omega_p \to [0,1]$ denote the CDF for pricing decision during period t under the optimal heuristic π^*. Also, let $\bar{r}_j^t$ and $\bar{\lambda}_j^t$ denote the expected revenue and demand rate from location j during period t under π^* (since we only consider the single-product setting, there is no need to use subscript k), i.e.,

$$\bar{r}_j^t := \mathbb{E}^{\pi^*}[R_j^t(p^t)] = \int_{\Omega_p} r_j(p)\ dF^t(p) \quad \text{and} \quad \bar{\lambda}_j^t := \mathbb{E}^{\pi^*}[D_j^t(p^t)] = \int_{\Omega_p} \lambda_j(p)\ dF^t(p).$$

To prove Lemma 1, we first show that there exist weight vectors $\{\alpha^t\}$ such that, for the uniform grid q^*defined in Section 4 and some sufficiently small $\epsilon_r, \epsilon_\lambda > 0$, the following hold:

$$\left|\bar{r}_j^t - \sum_{m=1}^M \alpha_m^t r_j(q_m^*)\right| = \left|\int_{\Omega_p} p\lambda_j(p)\ dF^t(p) - \sum_{m=1}^M \alpha_m r_j(q_m^*)\right| \le \epsilon_r \quad \forall j,t, \tag{EC.1}$$

$$\left|\bar{\lambda}_j^t - \sum_{m=1}^M \alpha_m^t \lambda_j(q_m^*)\right| = \left|\int_{\Omega_p} \lambda_j(p)\ dF^t(p) - \sum_{m=1}^M \alpha_m^t \lambda_j(q_m^*)\right| \le \epsilon_\lambda \quad \forall j,t, \tag{EC.2}$$

$$\sum_{m=1}^M \alpha_m^t = 1, \qquad \alpha_m^t \ge 0, \quad \forall m,t. \tag{EC.3}$$

Define a uniform partition of the interval Ω_p as $\Omega_p = \cup_{m=1}^M \mathcal{P}_m := \left[\cup_{m=1}^{M-1} [p_l + (m-1)\Delta_q, p_l + m\Delta_q)\right] \cup [p_u - \Delta_q, p_u]$, where $\Delta_q := (p_u - p_l)/M$ is the length of the sub-intervals. Then q^* can be expressed explictly as $q^* = (p_l + (m-1/2)\Delta_q)_{m=1}^M$. Consider a choice of weight vector $\alpha_m^t = \int_{\mathcal{P}_m} dF^t(p)$. Note that (EC.3) is satisfied immediately by definition. We now show that the combination of q^* and α^t defined above satisfy (EC.1) and (EC.2). By definition, for all $j \in [J]$, we have

$$\begin{aligned}\left|\bar{\lambda}_j^t - \sum_{m=1}^M \alpha_m^t \lambda_j(q_m^*)\right| &= \left|\int_{\Omega_p} \lambda_j(p)\ dF^t(p) - \sum_{m=1}^M \alpha_m^t \lambda_j(q_m^*)\right| = \left|\sum_{m=1}^M \int_{\mathcal{P}_m} (\lambda_j(p) - \lambda_j(q_m^*))\ dF^t(p)\right| \\ &\le \sum_{m=1}^M \int_{\mathcal{P}_m} |\lambda_j(p) - \lambda_j(q_m^*)|\ dF^t(p) \le \lambda_u \Delta_q.\end{aligned}$$

where the first inequality follows from triangular inequality and the last inequality follows from Assumption A1 together with $\lambda_u := \max_{j\in[J], p\in\Omega_p} |\lambda_j'(p)|$. By similar argument, since $|r_j'(p)| \le |\lambda_j(p) + p\lambda_j'(p)| \le 1 + p_u\lambda_u$ for all $p \in \Omega_p$, it is not difficult to show that (EC.1) is satisfied for $\epsilon_r = (1 + p_u\lambda_u)\Delta_q$.

We now show that the choices of q^* and α^t above guarantees a good approximation. Let $\tilde{\mathbf{x}}^t$ denote the optimal solution to the following LP:

$$\mathbf{FC}^A := \min_{\{x_{ij}^t\}} \left\{ \sum_{t=1}^T \sum_{i=0}^I \sum_{j=1}^J c_{ij} x_{ij}^t\ :\ \sum_{i=0}^I x_{ij}^t = \sum_{m=1}^M \alpha_m^t \lambda_j(q_m^*),\ \sum_{t=1}^T \sum_{j=1}^J x_{ij}^t \le C_i,\ 0 \le x_{ij}^t \le 1 \right\}.$$

Also, let $\bar{\mathbf{x}}^t$ denote the optimal solution of the following LP:

$$\mathbf{FC}^O := \min_{\{x_{ij}^t\}} \left\{ \sum_{t=1}^T \sum_{i=0}^I \sum_{j=1}^J c_{ij} x_{ij}^t\ :\ \sum_{i=0}^I x_{ij}^t = \bar{\lambda}_j^t,\ \sum_{t=1}^T \sum_{j=1}^J x_{ij}^t \le C_i,\ 0 \le x_{ij}^t \le 1 \right\}.$$

The only difference between $\mathbf{FC}^A$ and $\mathbf{FC}^O$ is on the RHS of fulfillment constraint. Since the optimal value of an LP is convex and thus Lipschitz continuous on a compact set in its RHS parameters (see Theorem 5.1 in Bertsimas and Tsitsiklis 1997) there exists a constants $K>0$ such that $\mathbf{FC}^A - \mathbf{FC}^O \le K\bar{c}\lambda_u\Delta_q$, where $\bar{c}=\max_{i\in[I]\cup\{0\},j\in[J]} c_{ij}$. So,

$$
\begin{aligned}
&\mathcal{J}^* - \mathcal{J}^{ALP}\\
&\le \left[\sum_{t=1}^{T}\sum_{j=1}^{J}\bar{r}_j^t - \mathbf{FC}^O\right] - \left[\sum_{t=1}^{T}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^t r_j(q_m^*) - \mathbf{FC}^A\right] \le \sum_{t=1}^{T}\sum_{j=1}^{J}\left|\bar{r}_j^t - \sum_{m=1}^{M}\alpha_m^t r_j(q_m^*)\right| + \left(\mathbf{FC}^A - \mathbf{FC}^O\right)\\
&\le \left[TJ(1+p_u\lambda_u)+K\bar{c}\lambda_u\right]\Delta_q \le \frac{(p_u-p_l)\left[TJ(1+p_u\lambda_u)+K\bar{c}\lambda_u\right]}{M}.
\end{aligned}
$$

The proof is concluded by letting $M=\lceil (p_u-p_l)\left[TJ(1+p_u\lambda_u)+K\bar{c}\lambda_u\right]/\epsilon\rceil+1$. $\square$

EC.2. Proof of Theorem 1

Let q^* be the set of discrete prices defined in Section 4. Without loss of generality, we assume that $T=1$. We consider a variant of RPF (V-RPF) defined as follow: during period t, fulfill the order from location j according to $\sigma_k^t(j)$ regardless of the availability of the corresponding FC; if the FC runs out of inventory, the retailer incurs a penalty cost of $\bar{c}:=2\cdot \max_{j\in[J],k[K]} c_{0jk}$. In other words, V-RPF incurs the same revenue as RPF, yet no smaller fulfillment cost. Consequently, the loss can be bounded as follows:

$$
\begin{aligned}
&\mathcal{J}^{ALP}(\theta) - \mathbb{E}\left[R^{RPF}(\theta)\right] \le \mathcal{J}^{ALP}(\theta) - \mathbb{E}\left[R^{V-RPF}(\theta)\right]\\
&= \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\theta}\sum_{j=1}^{J}(p^t)^\top D_j^t(p^t)\right] + \bar{c}\,\mathbb{E}\left[\sum_{i=1}^{I}\sum_{k=1}^{K}\left(\sum_{t=1}^{\theta}\sum_{j=1}^{J}X_{ijk}^t - C_{ik}(\theta)\right)^+\right]\\
&\qquad + \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}X_{ijk}^t - \sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}x_{ijk}^*\right]\\
&= \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{j=1}^{J}\Delta R_j^t\right] + \bar{c}\,\mathbb{E}\left[\sum_{i=0}^{I}\sum_{k=1}^{K}\left(\sum_{t=1}^{\theta}\sum_{j=1}^{J}X_{ijk}^t - C_{ik}(\theta)\right)^+\right] + \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}\Delta X_{ijk}^t\right],
\end{aligned}
$$

where $\Delta R_j^t := \sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - (p^t)^\top D_j^t(p^t)$, and $\Delta X_{ijk}^t := X_{ijk}^t - x_{ijk}^*$. By definition of RPF, $\mathbb{E}[\Delta R_j^t] = \mathbb{E}[\Delta X_{ijk}^t] = 0$. As for the last term, by triangular inequality,

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{i=0}^{I}\sum_{k=1}^{K}\left(\sum_{t=1}^{\theta}\sum_{j=1}^{J}X_{ijk}^t - C_{ik}(\theta)\right)^+\right]\\
&\le \mathbb{E}\left[\sum_{i=0}^{I}\sum_{k=1}^{K}\left(\sum_{t=1}^{\theta}\sum_{j=1}^{J}X_{ijk}^t - \theta\sum_{j=1}^{J}x_{ijk}^*\right)^+\right] + \mathbb{E}\left[\sum_{i=0}^{I}\sum_{k=1}^{K}\left(\theta\sum_{j=1}^{J}x_{ijk}^* - C_i(\theta)\right)^+\right]\\
&\le \sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}\mathbb{E}\left[\left(\sum_{t=1}^{\theta}X_{ijk}^t - x_{ijk}^*\right)^+\right] + 0 \le \sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}\left[\mathbf{Var}\left(\sum_{t=1}^{\theta}\Delta X_{ijk}^t\right)\right]^{1/2} = O\left(\sqrt{\theta}\right),
\end{aligned}
$$

where the second inequality follows from the inventory constraint in ALP, the last inequality follows because ΔX_{ijk}^t's are independent and bounded from above by $D_{jk}^t \le 1$. This completes the proof. $\square$

EC.3. Proof of Theorem 2

Let $T=1$. Per our discussion in Section 6, we can assume $\sum_{j=1}^{J} x^*_{ijk} = C_{ik}$ without loss of generality. Let $C_i^t(\theta)$ be the on-hand inventory level in FC i at the beginning of period t for a problem with size θ. By definition, we have $C_i^1(\theta)=\theta C_i$. Fix $\theta>0$. We divide our proof into several steps.

Step 1

We first define additional terms that will be useful for the proof:

$$\Delta^t_{jk} := D^t_{jk}(p^t) - \lambda_{jk}(p^t), \quad \tilde{\Delta}^t_{jk} := D^t_{jk}(p^t) - \sum_{m=1}^{M} \alpha^*_m \lambda_{jk}(q^t_m), \ \text{ and } \ \Delta y^t_{ijk} := Y^t_{ijk} - y^t_{ijk}.$$

To be precise, all the above terms should be written as a function of p^t and q^t. We suppress these notational dependencies for brevity. The term Δ^t_{jk} can be interpreted as the size of demand randomness induced by p^t; the term $\tilde{\Delta}^t_{jk}$ as the size of randomness induced by the sampling procedure; the term Δy^t_{ijk} as the size of randomness in the fulfillment outcome. Together with the term ΔC^t_{ik} defined in Section 6, all the above deltas are bounded random variables with zero mean for all $t \le \theta$. The proof is not difficult and is omitted.

Below we discuss two key observations that are useful to help us express the evolution of fulfillment decisions over time. We call an FLP^t to be "balanced" if its input parameters (q^t, C^t) satisfy (i) $\sum_{j=1}^{J}\sum_{m=1}^{M} \alpha^*_m \lambda_{jk}(q^t_m) = \sum_{i=0}^{I} C^t_{ik}/(T-t+1)$ for all k, and (ii) $C^t_{ik}>0$ for all i,k. We make our first observation regarding the solution of a balanced FLP^t.

Observation EC.1. *The optimal solution x^t to a non-DR-degenerate balanced FLP$^t(q^t, C^t)$ has the following property: For every $k\in[K]$, there are exactly $I+J$ strictly positive components in $(x^t_{ijk})_{i\in\{0\}\cup[I], j\in[J]}$, with the other components equal to zero. Moreover, the inventory constraints are all binding.*

Proof. Note that FLP$^t(q^t, C^t)$ is separable over k, so solving FLP$^t(q^t, C^t)$ is equivalent to solving K sub-problems defined below:

$$\mathrm{FLP}^t_k(q^t, C^t_k) := \left\{ \min_{x_{ijk}\ge 0} \sum_{i=0}^{I}\sum_{j=1}^{J} c_{ijk}x_{ijk} : \sum_{i=0}^{I} x_{ijk} = \sum_{m=1}^{M} \alpha^*_m \lambda_{jk}(q^t_m), \sum_{j=1}^{J} x_{ijk} \le C^t_{ik}/(T-t+1) \right\}.$$

Since FLP$^t(q^t, C^t)$ is balanced, all the inventory constraints in FLP$^t_k(q^t, C^t_k)$ must be binding. Since FLP$^t(q^t, C^t)$ is non-DR-degenerate and separable over k, FLP$^t_k(q^t, C^t_k)$ is also non-degenerate for each k. Thus, Observation EC.1 follows directly from the standard result on transportation LP (see Corollary 7.2 in Dantzig and Thapa 2006). □

Let $\mathbf{x}_k = (x_{ijk})_{i\in\{0\}\cup[I], j\in[J]}$ and $\mathbf{c}_k = (c_{ijk})_{i\in\{0\}\cup[I], j\in[J]}$. Given our assumptions in the statement of Theorem 2 and at the beginning of this section, FLP$^1(q^*, C)$ is non-DR-degenerate and balanced. Thus, for all k, FLP$^1_k(q^*, C_k)$ are non-degenerate and has $I+J$ non-zero components in $\mathbf{x}^*_k$ (since there are $I+J+1$ constraints with exactly one redundant). Let A_k and Q_k denote the coefficient matrix and the RHS of inventory constraints in FLP^1_k. Let $\bar{A}_k$ be the matrix where we delete the $(J+1)^{th}$ row from A_k, i.e., the row corresponding to the inventory constraint on FC 0, and $\bar{Q}_k$ be the vector where we delete C_{0k}/θ from Q_k. This constraint is redundant, since any $\mathbf{x}_k$ satisfying the system of equations $\bar{A}_k\mathbf{x}_k = \bar{Q}_k$ automatically satisfies $\sum_{j=1}^{J} x^t_{0jk} = C_{0k}/\theta$ (the deleted constraint). By Theorem 2.5 in Bertsimas and Tsitsiklis (1997), the FLP^1_k

is equivalent to $\{\min\ \mathbf{c}_k^\top \mathbf{x}_k:\ \bar{A}_k\mathbf{x}_k = \bar{Q}_k,\ \mathbf{x} \succeq \mathbf{0}\}$; moreover, by Lemma 7.1 in Dantzig and Thapa (2006), $\bar{A}_k$ has linearly independent rows. Let $\mathcal{B}_k = \{(i,j):\ 0 < x^*_{ijk} < 1\}$ and $\mathcal{N}_k = \{(i,j):\ x^*_{ijk} = 0\}$ be the indices of the optimal basic and non-basic variables respectively. Without loss of generality, we assume that $\bar{A}_k$ is written as $[B_k, N_k]$ where B_k and N_k are the sub-matrices of $\bar{A}_k$ corresponding to the basic and non-basic indices in $\mathcal{B}_k$ and $\mathcal{N}_k$ respectively. Following the same decomposition, the optimal solution can be represented as $\mathbf{x}^*_k = [\mathbf{x}^*_{k,B}, \mathbf{x}^*_{k,N}]$, where $\mathbf{x}^*_{k,B} = B_k^{-1}\bar{Q}_k$ and $\mathbf{x}^*_{k,N} = \mathbf{0}$ (the invertibility of B_k is proved in Theorem 7.6 in Dantzig and Thapa 2006). Thus, the unique optimal solution to FLP^1 can be accordingly written as $\mathbf{x}^* = [\mathbf{x}^*_B; \mathbf{x}^*_N]$, where $\mathbf{x}^*_B = (\mathbf{x}^*_{k,B})_{k=1}^K$, $\mathbf{x}^*_N = (\mathbf{x}^*_{k,N})_{k=1}^K$. Note that if we define $B = \mathrm{diag}(B_1, \ldots, B_K)$ as a block diagonal matrix with $(B_k)_{k=1}^K$ as its main diagonal blocks and zero matrices as off-diagonal blocks, and $\bar{Q} = [\bar{Q}_1; \ldots; \bar{Q}_K]$, we can write $\mathbf{x}^*_B = B^{-1}\bar{Q}$. Let Q_k^t be the RHS of FLP$_k^t$ and $\bar{Q}_k^t$ be the vector where we delete $C_{0k}^t/(\theta - s)$ from Q_t^k. Define $\delta Q_k^t := ((\sum_{m=1}^M \alpha_m^*\ \lambda_{jk}(q_m^t) - \sum_{m=1}^M \alpha_m^* \lambda_{jk}(q_m^*))_{j=1}^J, (-\sum_{s=1}^{t-1} \Delta C_{ik}^s/(\theta - s))_{i=0}^I)$ and let $\delta\bar{Q}_k^t$ be the vector where we delete $-\sum_{s=1}^{t-1}\Delta C_{0k}^s/(\theta - s)$ from δQ_k^t. Let $\delta\bar{Q}^t = (\delta\bar{Q}_k^t)_{k=1}^K$. Following the same decomposition, we will also write $\mathbf{c} = [\mathbf{c}_B; \mathbf{c}_N]$. Per our definition in Section 3, $\lambda^{tot}(p)$ is the aggregated purchase probability given a price vector $p \in \Omega_p$. We make our second observation below:

Observation EC.2. *At period t, as long as the following conditions hold:*

$$\sum_{j=1}^J \lambda_j(q_m^t) = \hat{\lambda}_m^t := \lambda^{tot}(q_m^*) - \frac{1}{M\alpha_m^*}\left(\sum_{i=0}^I \sum_{s=1}^{t-1} \frac{\Delta C_i^s}{T - s}\right) \in \otimes_{k=1}^K [0,1], \tag{EC.4}$$

$$C_{ik}^t(\theta) = \hat{C}_{ik}^t(\theta) := (\theta - t + 1)\left[C_{ik} - \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s}\right] \geq 0, \tag{EC.5}$$

$$\mathbf{x}^*_{k,B} + B_k^{-1}\delta\bar{Q}_k^t \succeq \mathbf{0}, \tag{EC.6}$$

then the unique optimal solution to FLP^t is given by $\mathbf{x}^t_{k,B} = \mathbf{x}^*_{k,B} + B_k^{-1}\delta\bar{Q}_k^t$ and $\mathbf{x}^t_{k,N} = \mathbf{0}$ for all k.

Proof. Under condition (EC.4), FLP^t is balanced. This is so because, for all k,

$$\sum_{j=1}^J \sum_{m=1}^M \alpha_m^* \lambda_{jk}(q_m^t) = \sum_{j=1}^J \sum_{m=1}^M \alpha_m^* \lambda_{jk}(q_m^*) - \sum_{i=0}^I \sum_{s=1}^t \frac{\Delta C_{ik}^s}{T - s} = \sum_{i=0}^I C_{ik} - \sum_{i=0}^I \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s} = \sum_{i=0}^I \frac{C_{ik}^t(\theta)}{\theta - t + 1},$$

where the second equality follows from our assumption in the beginning of this section, and the last equality follows from the definition of ΔC_{ik}^t. As a result, for all k, the inventory constraints in FLP$_k^t$ are all binding. Notice that condition (EC.4) and (EC.5) implies that $Q_k^t = Q_k + \delta Q_k^t \succeq \mathbf{0}$, and thus FLP$_k^t$ is equivalent to $\left\{\min_{\mathbf{x}_k^t}\ \mathbf{c}_k^\top \mathbf{x}_k^t : \bar{A}_k \mathbf{x}_k^t = \bar{Q}_k + \delta\bar{Q}_k^t,\ \mathbf{x}_k^t \succeq \mathbf{0}\right\}$. The feasibility of the proposed optimal solution can be directly verified under condition (EC.6); its optimality follows from Karush-Kuhn-Tucker (KKT) conditions; and its uniqueness follows from the invertibility of B_k. $\square$

Step 2

Define $\hat{\mathbf{x}}^t := (\hat{\mathbf{x}}_B^t, \mathbf{x}_N) = (\mathbf{x}^*_B + B^{-1}\delta\bar{Q}^t, \mathbf{0})$. Let $\phi_x = \min_{k\in[K]} \min_{(i,j)\in\mathcal{B}_k} x^*_{ijk} > 0$ (by non-degeneracy assumption); $\Phi_1 = \max_{p\in\Omega_p,\ j\in[J],\ k,\ell\in[K]} |\partial\lambda_{jk}(p)/\partial p_\ell| > 0$ (it is finite by Assumption A1); $\Phi_2 = \max_{k\in[K]} ||B_k^{-1}||_\infty > 0$ (it is also finite by the invertibility of B_k); $\phi_\lambda := \max\{x > 0 : \lambda^{tot}(q_m^*) + x \cdot \mathbf{1} \in \otimes_{k=1}^K [0,1],\ \forall m\} > 0$ (by Assumption A1 and the fact that q_m^* lies in the interior of Ω_p); and $v > 0$ denote the smallest absolute eigenvalue of $\mathbf{J}_{\lambda^{tot}}$ (by Assumption A3). Remember that, we assumed without loss of generality that $\alpha^* \succ \mathbf{0}$ since we can delete any α_m^* with zero value without changing anything else. We state a lemma.

LEMMA EC.1. *Suppose that $\lambda^{tot}(q_m^s) = \hat{\lambda}_m^s \in \otimes_{k=1}^K [0,1]$, $\mathbf{x}_s = \hat{\mathbf{x}}_s \succeq \mathbf{0}$ and $C_i^s(\theta) = \hat{C}_i^s(\theta) \succeq \mathbf{0}$ for all $s < t$. Then $\lambda^{tot}(q_m^t) = \hat{\lambda}_m^t$, $\mathbf{x}^t = \hat{\mathbf{x}}^t$ and $C_i^t(\theta) = \hat{C}_i^t(\theta)$ hold if the following two conditions hold at time t*

$$(\dagger): \qquad \left| \sum_{i=1}^{I} \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s} \right| \le \min\left\{ \frac{\phi_x}{\Phi_2} \left(1 + \frac{K\Phi_1}{v}\right)^{-1},\ \phi_\lambda M \cdot \min_{m\in[M]} \alpha_m^* \right\},\ \forall k,$$

$$(\dagger\dagger): \qquad \left| \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s} \right| \le C_{ik},\ \forall i, k,$$

Proof. We proceed by induction. The base case ($t = 1$) can be verified directly by definition. Now, consider $t > 1$. Assume the identity holds for $s \le t-1$. Given condition (†) and the definition of ϕ_λ, it is not difficult to show that $\hat{\lambda}_m^t \in \otimes_{k=1}^K [0,1]$. Since $\lambda^{tot}(q_m^t)$ is simply the projection of $\hat{\lambda}_m^t$ onto $\otimes_{k=1}^K [0,1]$ (see Step 2a in R^{2}PF), $\lambda^{tot}(q_m^t) = \hat{\lambda}_m^t$.

We now show that $C_{ik}^t(\theta) = \hat{C}_{ik}^t(\theta)$. Suppose that, in Step 2c of R^{2}PF, we sample m^t for some $m^t \in [M]$. Remember that, in period $t-1$, the probability of using FC i to fulfill the request of product k from location j conditioned on $D_{jk} = 1$ is $y_{ijk}^{t-1} = x_{ijk}^{t-1} / \sum_{i=0}^I x_{ijk}^{t-1}$. Moreover, since conditions (EC.4) - (EC.6) are implied for all $s \le t$ by the inductive assumption, by Observation EC.2, the inventory constraints in FLP^{t-1} are binding. So, the remaining inventory at the beginning of period t satisfies:

$$\begin{aligned} C_{ik}^t(\theta) &= C_{ik}^{t-1}(\theta) - \sum_{j=1}^J X_{ijk}^{t-1} = C_{ik}^{t-1}(\theta) - \sum_{j=1}^J y_{ijk}^{t-1} \left(\sum_{m=1}^M \alpha_m^* \lambda_{jk}(q_m^{t-1}) \right) - \Delta C_{ik}^{t-1} \\ &= C_{ik}^{t-1}(\theta) - \sum_{j=1}^J x_{ijk}^{t-1} - \Delta C_{ik}^{t-1} = C_{ik}^{t-1}(\theta) - \frac{C_{ik}^{t-1}(\theta)}{\theta - t + 2} - \Delta C_{ik}^{t-1} \\ &= (\theta - t + 2 - 1) \left[C_{ik}(\theta) - \sum_{s=1}^{t-2} \frac{\Delta C_{ik}^s}{\theta - s} \right] - \Delta C_{ik}^{t-1} = \hat{C}_i^t(\theta), \end{aligned}$$

where the second equality follows from the definition of ΔC_{ik}^t; the third equality follows from the fulfillment constraint in FLP^t; the fourth constraint follows since the inventory constraints in FLP^{t-1} are binding; and, the fifth constraints follows from the inductive assumption.

At last, to show that $\mathbf{x}^t = \hat{\mathbf{x}}^t$, by Observation EC.2, it suffices to show conditions (EC.4) - (EC.6) are satisfied for period t. Condition (EC.4) is implied by $\lambda^{tot}(q_m^t) = \hat{\lambda}_m^t$. Since condition (††) implies $\hat{C}_{ik}^t(\theta) \ge 0$, and we have shown that $C_{ik}^t(\theta) = \hat{C}_{ik}^t(\theta)$, condition (EC.5) is satisfied. To check condition (EC.6), define $\delta q_m^t = q_m^t - q_m^*$. By Assumption A1 and Mean Value Theorem, $\delta q_m^t = [\mathbf{J}_{\lambda^{tot}}(\xi_m^t)]^{-1} \left(\sum_{i=0}^I \sum_{s=1}^{t-1} \Delta C_i^s / (\theta - s) \right) / (M\alpha_m^*)$ for some $\xi_m^t \in \Omega_p$. By Mean Value Theorem again, there exist $\zeta_{mk}^t \in \Omega_p$ such that

$$\left| \sum_{m=1}^M \alpha_m^* \left[\lambda_{jk}(q_m^t) - \lambda_{jk}(q_m^*) \right] \right| = \left| \sum_{m=1}^M \frac{(\nabla \lambda_{jk}(\zeta_{mk}^t))^\top [\mathbf{J}_{\lambda^{tot}}(\xi_m^t)]^{-1}}{M} \left(\sum_{i=0}^I \sum_{s=1}^{t-1} \frac{\Delta C_i^s}{\theta - s} \right) \right| \le \frac{K\Phi_1}{v} \max_{k\in[K]} \left| \sum_{i=0}^I \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s} \right|$$

where the inequality holds by Assumption A3 and the definition of Φ_1. So,

$$\left|\left| B_k^{-1} \delta \bar{Q}_k^t \right|\right|_\infty \le \left|\left| B_k^{-1} \right|\right| \cdot \left|\left| \delta \bar{Q}_k^t \right|\right| \le \Phi_2 \cdot \left(1 + \frac{K\Phi_1}{v}\right) \max_{k\in[K]} \left| \sum_{i=0}^I \sum_{s=1}^{t-1} \frac{\Delta C_{ik}^s}{\theta - s} \right| \le \phi_x,$$

where the last inequality follows from condition (†). This implies condition (EC.6). □

Step 3

In this step, we show that the conditions in Lemma EC.1 hold for the majority of the selling season. Define a stopping time $\tau(\theta)$ to be the first t such that either (†) or (††) is violated. According to Lemma EC.1, for any period before $\tau(\theta)$, we can explicitly characterize the evolution of price, fulfillment assignment, and inventory consumption. The following lemma provides a lower bound on the length of $\tau(\theta)$.

LEMMA EC.2. *There exists a constant $\Psi_3 > 0$ independent of θ such that*

$$\mathbb{E}\left[\theta - \tau(\theta)\right] \le \Psi_3(1 + \log\theta).$$

Proof. Define $\tau_1(\theta)$ and $\tau_2(\theta)$ to be the first period t such that conditions (†) and (††) are violated, respectively. By definition $\tau(\theta) = \min_{i\in\{1,2\}} \tau_i(\theta)$. In what follows, we will only bound $\tau_1(\theta)$, since $\tau_2(\theta)$ can be bounded using a similar argument.

Let Γ_k denote the RHS of the inequality in condition (†) in Lemma EC.1. The sequence

$$\left\{ S_k^t = \sum_{i=0}^{I} \frac{\Delta C_{ik}^{t-1}}{\theta-(t-1)} + \sum_{i=0}^{I} \frac{\Delta C_{ik}^{t-2}}{\theta-(t-2)} + \cdots + \sum_{i=0}^{I} \frac{\Delta C_{ik}^{1}}{\theta-1} \right\}_{t\le\theta}$$

is a Martingle with respect to the natural filtration $\{\mathcal{H}^t\}$, where $\mathcal{H}^t$ is the history of all information up to the beginning of period t. This implies that the sequence $\{|S_k^t|\}_{t\le\theta}$ is a sub-Martingle. By Doob's submartingle inequality (see for example Williams 1991) and union bound,

$$\mathbb{P}(\tau_1(\theta)\le t) \;\le\; \mathbb{P}\left(|S_k^s| \ge \Gamma_k \text{ for some } s\le t, k\in[K]\right) \le \sum_{k=1}^{K} \mathbb{P}\left(\max_{s\le t} |S_k^s| \ge \Gamma_k\right) \;\le\; \sum_{k=1}^{K} \frac{\mathbb{E}\left[(S_k^t)^2\right]}{\Gamma_k^2}.$$

Note that ΔC_{ik}^s and ΔC_{jk}^t are independent for all $s\ne t$ and $i,j\in\{0\}\cup I$. So,

$$\mathbb{E}[(S_k^t)^2] = \mathbb{E}\left[\left(\sum_{s=1}^{t-1}\sum_{i=0}^{I}\frac{\Delta C_{ik}^s}{\theta-s}\right)^2\right] = \sum_{s=1}^{t-1}\frac{\mathbb{E}\left[\left(\sum_{i=0}^{I}\Delta C_{ik}^s\right)^2\right]}{(\theta-s)^2} = \sum_{s=1}^{t-1}\frac{\sum_{i,j\in\{0\}\cup[I]}\mathbb{E}\left[\Delta C_{ik}^s \Delta C_{jk}^s\right]}{(\theta-s)^2} \;=\; O\left(\frac{1}{\theta-t}\right),$$

where the last inequality follows from the boundedness of $\mathbb{E}\left[\Delta C_{ik}^s \Delta C_{jk}^s\right]$. The proof is complete by noting that $\mathbb{E}\left[\theta-\tau_1(\theta)\right] \;=\; \sum_{t=2}^{\theta}\mathbb{P}(\tau_1(\theta)\le t) \;=\; 1+\sum_{t=2}^{\theta-1} O\left(\frac{1}{\theta-t}\right) \;=\; O(\log\theta)$. □

Step 4

We now bound the loss of R^2PF. First, note that we can decouple the loss into two terms as follows:

$$\begin{aligned}
&\mathcal{J}^{LP}(\theta) - \mathbb{E}\left[R^{R^2PF}(\theta)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\theta}\sum_{j=1}^{J}(p^t)^\top D_j^t(p^t)\right] + \mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}X_{ijk}^t - \sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}x_{ijk}^*\right].
\end{aligned}$$

The two terms on the RHS of the equation above are the loss in revenue and the loss in fulfillment cost of R^2PF, respectively. We start with providing an upper bound for the loss in revenue:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\theta}\sum_{j=1}^{J}(p^t)^\top D_j^t(p^t)\right] \\
&\le \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J} R_j^t(p^t)\right] + \mathbb{E}\left[\sum_{t=\tau(\theta)}^{\theta}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*)\right]
\end{aligned}$$

$$
\begin{aligned}
&= \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J} R_j^t(p^t)\right] + \mathbb{E}\left[(\theta-\tau(\theta)+1)\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*)\right] \\
&\le \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J} R_j^t(p^t)\right] + Kp_u(1+\Psi_3+\Psi_3\log\theta), \qquad \text{(EC.7)}
\end{aligned}
$$

where the last inequality follows from Lemma EC.2, the boundedness of price, and the assumption of at most one arrival per period. Let $\hat{\Delta}_j^t = \sum_{m=1}^{M}\alpha_m^* r_j(q_m^t) - (p^t)^\top D_j(p^t)$. Define $r^{tot}(p) = \sum_{j=1}^{J} r_j(p) = p^\top \lambda^{tot}(p)$. By Assumption A1, there exists an inverse of $\lambda^{tot}(p)$, which we will denote as $p(\lambda^{tot}) : \otimes_{k=1}^{K}[0,1] \to \Omega_p$. With slight abuse of notation, we will use $r^{tot}(\lambda^{tot}) = (p(\lambda^{tot}))^\top \lambda^{tot}$ to denote total revenue rate as a function of aggregate demand. Let $\lambda_m^* = \lambda^{tot}(q_m^*)$, $\lambda_m^t = \lambda^{tot}(q_m^t)$, and $\epsilon^t = \sum_{i=0}^{I}\sum_{s=1}^{t-1}\Delta C_i^s/(\theta-s)$. For $t \le \tau(\theta)$, we know that$\lambda_m^t = \lambda_m^* - \epsilon^t/(M\alpha_m^*)$. By Taylor's expansion at λ_m^*, we have

$$
\begin{aligned}
r^{tot}(q_m^t) = r^{tot}(\lambda_m^t) &= r^{tot}(\lambda_m^*) - (\nabla r^{tot}(\lambda_m^*))^\top \epsilon_m^t/(M\alpha_m^*) + (\epsilon^t)^\top \nabla^2 r^{tot}(\eta^t)\epsilon^t/(2M^2(\alpha_m^*)^2) \\
&= r^{tot}(q_m^*) - (\nabla r^{tot}(\lambda_m^*))^\top \epsilon_m^t/(M\alpha_m^*) + (\epsilon^t)^\top \nabla^2 r^{tot}(\eta^t)\epsilon^t/(2M^2(\alpha_m^*)^2)
\end{aligned}
$$

for some $\eta_m^t \in \otimes_{k=1}^{K}[0,1] \in \Omega_p$. So, the first term in (EC.7) can be bounded as follows:

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^*) - \sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J} R_j^t(p^t)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{m=1}^{M}\alpha_m^* r^{tot}(q_m^*) - \sum_{t=1}^{\tau(\theta)-1}\sum_{m=1}^{M}\alpha_m^* r^{tot}(q_m^t)\right] + \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{m=1}^{M}\alpha_m^* r_j(q_m^t) - \sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}(p^t)^\top D_j(p^t)\right] \\
&\le \mathbb{E}\left[\sum_{t=2}^{\tau(\theta)-1}\sum_{m=1}^{M}\frac{(\nabla r^{tot}(\lambda_m^*))^\top \epsilon^t}{M}\right] - \mathbb{E}\left[\sum_{t=2}^{\tau(\theta)-1}\sum_{m=1}^{M}\frac{(\epsilon^t)^\top \nabla^2 r^{tot}(\eta_m^t)\epsilon^t}{2M^2\min_{m\in[M]}\alpha_m^*}\right] + \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)}\sum_{j=1}^{J}\hat{\Delta}_j^t\right] + Kp_u, \qquad \text{(EC.8)}
\end{aligned}
$$

where the last inequality holds because $\mathbb{E}[\hat{\Delta}_j^{\tau(\theta)}] \le Kp_u$. Note that $\{\sum_{j=1}^{J}\hat{\Delta}_j^t\}_{t\le\theta}$ is a Martingale with respect to $\{\mathcal{H}^t\}_{t\le\theta}$ and $\tau(\theta)$ is bounded. So, by stopping time theorem (Williams 1991), $\mathbb{E}[\sum_{t=1}^{\tau(\theta)}\sum_{j=1}^{J}\hat{\Delta}_j^t] = 0$. We are left to bound the first two terms in (EC.8). Note that $\mathbb{E}[\sum_{t=2}^{\tau(\theta)-1}\epsilon^t] = \mathbb{E}[\sum_{t=2}^{\tau(\theta)}\epsilon^t] - \mathbb{E}[\sum_{t=\tau(\theta)}^{\theta}\epsilon^t] = -\mathbb{E}[\sum_{t=\tau(\theta)}^{\theta}\epsilon^t]$. By stopping time theorem again, $\mathbb{E}[\epsilon^{\tau(\theta)}] = \mathbf{0}$, and $\mathbb{E}[\epsilon^t] = \mathbf{0}$ for all $t > \tau(\theta)$. Consequently, $\mathbb{E}[\sum_{t=2}^{\tau(\theta)-1}\sum_{m=1}^{M}(\nabla r^{tot}(\lambda_m^*))^\top\epsilon^t] = (\sum_{m=1}^{M}\nabla r^{tot}(\lambda_m^*))^\top \mathbb{E}[\sum_{t=2}^{\tau(\theta)-1}\epsilon^t] = 0$. As for the second term in (EC.8), let $\Phi_3 > 0$ be the largest absolute eigenvalue of $\nabla^2 r^{tot}$. By Assumption A3, Φ_3 is finite. We thus have

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{t=2}^{\tau(\theta)-1}\sum_{m=1}^{M}(\epsilon^t)^\top \nabla^2 r^{tot}(\eta_m^t)\epsilon^t\right] \le \Phi_3 \mathbb{E}\left[\sum_{t=2}^{\tau(\theta)-1}\sum_{k=1}^{K}\left(\sum_{i=1}^{I}\sum_{s=1}^{t-1}\frac{\Delta C_{ik}^s}{\theta-s}\right)^2\right] \\
&= \Phi_3 \sum_{t=2}^{\tau(\theta)-1}\sum_{k=1}^{K}\sum_{1\le s,v\le t-1}\frac{\mathbb{E}\left[\left(\sum_{i=1}^{I}\Delta C_{ik}^s\right)^2\left(\sum_{i=1}^{I}\Delta C_{ik}^v\right)^2\right]}{(\theta-s)(\theta-v)} = \Phi_3 \sum_{t=2}^{\tau(\theta)-1}\sum_{k=1}^{K}\sum_{s=1}^{t-1}\frac{\mathbb{E}\left[\left(\sum_{i=1}^{I}\Delta C_{ik}^s\right)^2\right]}{(\theta-s)^2} = O(\log\theta).
\end{aligned}
$$

At last we bound the loss of fulfillment cost. By Lemma EC.1, for $t < \tau(\theta)$, $\mathbf{x}^t = [x_B^* + B^{-1}\delta\bar{Q}^t; \mathbf{0}]$. By definition, $\bar{c}$ is larger than all unit shipping costs. So,

$$
\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}X_{ijk}^t - \sum_{t=1}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}x_{ijk}^*\right] \\
&\le \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}X_{ijk}^t - \sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}x_{ijk}^*\right] + \mathbb{E}\left[\sum_{t=\tau(\theta)}^{\theta}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K}c_{ijk}X_{ijk}^t\right]
\end{aligned}
$$

$$\leq \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}X_{ijk}^t - \sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}x_{ijk}^*\right] + \bar{c}IJK\mathbb{E}[\theta-\tau(\theta)+1]$$
$$\leq \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}X_{ijk}^t - \sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}x_{ijk}^*\right] + \bar{c}IJK(1+\Psi_3+\Psi_3\log\theta). \qquad \text{(EC.9)}$$

We are left to bound the first term in (EC.9). Let $\Delta x_{ijk}^t = X_{ijk}^t - x_{ijk}^t$. Since $\mathbf{x}^t = \hat{\mathbf{x}}^t$ for all $t<\tau(\theta)$, we have:

$$\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}X_{ijk}^t - \sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J} c_{ijk}x_{ijk}^*\right]$$
$$= \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{K} c_{ijk}\left(x_{ijk}^t - x_{ijk}^*\right)\right] + \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J} c_{ij}\Delta x_{ij}^t\right]$$
$$= \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1} \mathbf{c}_B^\top B^{-1}\delta\bar{Q}^t\right] + \mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J} c_{ij}\Delta x_{ij}^t\right]$$
$$\leq \bar{c}(I+J)K||B^{-1}||_1\left\{\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{j=1}^{J}\sum_{k=1}^{K}\sum_{m=1}^{M}\alpha_m^*\left(\lambda_{jk}(q_m^t)-\lambda_{jk}(q_m^*)\right) - \sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{k=1}^{K}\sum_{s=1}^{t}\frac{\Delta C_{ik}^s}{\theta-s}\right]\right\}$$
$$+\bar{c}\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J}\Delta x_{ij}^t\right]$$
$$= -2\bar{c}(I+J)K||B^{-1}||_1\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=0}^{I}\sum_{k=1}^{K}\sum_{s=1}^{t}\frac{\Delta C_{ik}^s}{\theta-s}\right] + \bar{c}\mathbb{E}\left[\sum_{t=1}^{\tau(\theta)-1}\sum_{i=1}^{I}\sum_{j=1}^{J}\Delta x_{ij}^t\right]$$

where the second inequality follows from the definition of $\delta\bar{Q}^t$, the second equality follows from the definition of $\tau(\theta)$ and Lemma EC.1. Note that $\{\sum_{s=1}^t \Delta x_{ij}^s\}_{t\leq\theta}$ is Martingale with respect to the filtration $\{\mathcal{H}^t\}_{t\leq\theta}$. Following a similar argument as in bounding the revenue loss, it is not difficult to see that the terms after the above equation can be bounded by a constant independent of θ. $\square$

EC.4. Parameters of Numerical Experiments 1

An arrival customer makes a purchase with probability $\exp(A_1+B_1p)$ with the following parameters:

$$A_1 = \begin{bmatrix}-1.477\\-1.300\\-3.419\\-4.418\\-4.951\\-1.542\\-3.467\\-3.035\\-1.103\end{bmatrix}, \quad B_1 = \begin{bmatrix}-6.300 & 0.333 & 0.300 & 0.233 & 0.283 & 0.250 & 0.183 & 0.283 & 0.300\\ 0.317 & -5.900 & 0.333 & 0.283 & 0.217 & 0.233 & 0.250 & 0.317 & 0.200\\ 0.183 & 0.333 & -6.400 & 0.183 & 0.333 & 0.283 & 0.333 & 0.333 & 0.333\\ 0.333 & 0.333 & 0.167 & -6.000 & 0.167 & 0.283 & 0.217 & 0.267 & 0.217\\ 0.267 & 0.250 & 0.317 & 0.167 & -5.400 & 0.300 & 0.267 & 0.183 & 0.200\\ 0.183 & 0.300 & 0.333 & 0.217 & 0.233 & -5.400 & 0.200 & 0.183 & 0.200\\ 0.217 & 0.183 & 0.283 & 0.167 & 0.300 & 0.283 & -5.700 & 0.200 & 0.267\\ 0.267 & 0.233 & 0.300 & 0.183 & 0.300 & 0.283 & 0.200 & -5.700 & 0.250\\ 0.333 & 0.333 & 0.300 & 0.317 & 0.200 & 0.183 & 0.250 & 0.200 & -5.800\end{bmatrix} \times 10^{-3}$$

The absolute magnitude of the components in A_j and B_j depends on the feasible price range, which, in our setting, depends on the shipping cost. The initial inventory is set as follows:

$$C = \begin{bmatrix}0.0196 & 0.0196 & 0.0196 & 0.0196 & 0.0196 & 0.0196 & 0.0196 & 0.0196 & 0.0196\\ 0.0130 & 0.0130 & 0.0130 & 0.0130 & 0.0130 & 0.0130 & 0.0130 & 0.0130 & 0.0130\\ 0.0122 & 0.0122 & 0.0122 & 0.0122 & 0.0122 & 0.0122 & 0.0122 & 0.0122 & 0.0122\\ 0.0189 & 0.0189 & 0.0189 & 0.0189 & 0.0189 & 0.0189 & 0.0189 & 0.0189 & 0.0189\\ 0.0127 & 0.0127 & 0.0127 & 0.0127 & 0.0127 & 0.0127 & 0.0127 & 0.0127 & 0.0127\\ 0.0136 & 0.0136 & 0.0136 & 0.0136 & 0.0136 & 0.0136 & 0.0136 & 0.0136 & 0.0136\end{bmatrix}$$