

Dynamic Pricing Without Knowing the Demand Function: Risk Bounds and Near-Optimal Algorithms

Omar Besbes

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104,
obesbes@wharton.upenn.edu

Assaf Zeevi

Graduate School of Business, Columbia University, New York, New York 10027,
assaf@gsb.columbia.edu

We consider a single-product revenue management problem where, given an initial inventory, the objective is to dynamically adjust prices over a finite sales horizon to maximize expected revenues. Realized demand is observed over time, but the underlying functional relationship between price and mean demand rate that governs these observations (otherwise known as the demand function or demand curve) is not known. We consider two instances of this problem: (i) a setting where the demand function is assumed to belong to a known parametric family with unknown parameter values; and (ii) a setting where the demand function is assumed to belong to a broad class of functions that need not admit any parametric representation. In each case we develop policies that learn the demand function “on the fly,” and optimize prices based on that. The performance of these algorithms is measured in terms of the *regret*: the revenue loss relative to the maximal revenues that can be extracted when the demand function is known prior to the start of the selling season. We derive lower bounds on the regret that hold for any admissible pricing policy, and then show that our proposed algorithms achieve a regret that is “close” to this lower bound. The magnitude of the regret can be interpreted as the economic value of prior knowledge on the demand function, manifested as the revenue loss due to model uncertainty.

Subject classifications: revenue management; pricing; estimation; learning; exploration-exploitation; value of information; asymptotic analysis.

Area of review: Stochastic Models.

History: Received June 2007; revisions received December 2007, June 2008; accepted June 2008. Published online in *Articles in Advance* April 24, 2009.

1. Introduction

Since the deregulation of the airline industry in the 1970s, revenue management practices that deal with complex pricing and demand management decisions have become increasingly prevalent in a variety of industries; see Bitran and Caldentey (2003), Elmaghraby and Keskinocak (2003), and Talluri and van Ryzin (2005) for a broad overview of the field and recent developments. (The literature review in §2 provides a detailed survey of work directly related to the current paper.) A critical assumption made in most academic studies of revenue management problems is that the functional relationship between the mean demand rate and price, often referred to as the *demand function* or *demand curve*, is known to the decision maker. This makes the underlying problems more tractable and allows one to extract structural insights. At the same time, this assumption of “full information” endows the decision maker with knowledge that s/he does not typically possess in practice.

Lack of information concerning the demand model raises several fundamental questions. First and foremost, is it possible to quantify the “value” of full information

(for example, by measuring the revenue loss due to imperfect information)? Second, is it possible to achieve anything close to the maximal revenues in the full-information setting by judiciously combining real-time demand learning and pricing strategies? Finally, how would such strategies exploit prior information, if any, on the structure of the demand function?

The main objective of this paper is to shed some light on the aforementioned questions. Our departure point will be a prototypical single-product revenue management problem first introduced and formalized by Gallego and van Ryzin (1994). This formulation models realized demand as a Poisson process whose intensity at each point in time is determined by a price set by the decision maker. Given an initial inventory, the objective is to dynamically price the product so as to maximize expected revenues over a finite selling horizon. In the dynamic optimization problem considered in Gallego and van Ryzin (1994), the decision maker knows the demand function prior to the start of the selling season and designs optimal policies based on this information. In the setting we pursue in this paper it is

only possible to observe *realized demand* over time, and the demand function itself is *not known*. To that end, we consider two levels of uncertainty with regard to the demand model: (i) a *nonparametric setting* where the demand function is only assumed to belong to a broad functional class satisfying mild regularity conditions; and (ii) a *parametric setting* in which the demand function admits a given parametric structure but the parameter values are not known.

The absence of perfect prior information concerning the demand model introduces an important new component into the above dynamic optimization problem, namely, tension between exploration (demand learning) and exploitation (pricing). The longer one spends learning the demand characteristics, the less time remains to exploit that knowledge and optimize profits. On the other hand, less time spent on demand learning leaves more residual uncertainty that could hamper any pricing strategy in the exploitation phase. One of the main contributions of this paper is to formulate this dynamic pricing problem under incomplete information, to pursue an analysis that highlights the key trade-offs discussed above, and to articulate them in a precise mathematical manner.

To address uncertainty with regard to the demand model, we introduce a family of pricing policies that learn the demand function “on the fly.” Their performance will be measured in terms of the revenue loss relative to a full-information benchmark that assumes knowledge of the demand function. We refer to this loss as the *regret* associated with not knowing the demand function a priori; the magnitude of the regret quantifies the *economic value* of prior model information. The policies we consider are designed with the objective of achieving a “small” regret *uniformly* over the relevant class of demand functions (either parametric or nonparametric). This adversarial setting, where nature is allowed to counter a chosen policy with the “worst” demand function, ensures that policies exhibit “good” performance irrespective of the true demand model.

The complexity of the problem described above makes it difficult to evaluate the performance of any reasonable policy, except via numerical experiments. To address this issue, we consider an asymptotic regime that is characterized by a high volume of sales. More specifically, the initial level of inventory and the magnitude of demand (“market size”) grow large in proportion to each other; see Gallego and van Ryzin (1994) and Talluri and van Ryzin (2005) for further examples in the revenue management literature that adopt this framework. This regime allows us to bound the magnitude of the regret and to establish a rather surprising result with regard to our proposed policies: as the sales volume grows large, the regret eventually shrinks to zero. That is, these policies achieve (asymptotically) the maximal full-information revenues despite the absence of prior information regarding the demand function; in that sense, they are *asymptotically optimal*.

In more detail, the main contributions of this paper are summarized as follows.

(i) We introduce a nonparametric pricing policy (see Algorithm 1) that requires almost no prior information on the demand function. In settings where the structure of the demand function is known up to the values of certain parameters, we develop a parametric pricing policy based on maximum-likelihood estimation (see Algorithm 2 and Algorithm 3).

(ii) We establish lower bounds on the regret that hold for *any* admissible learning and pricing policy (see Propositions 2 and 4).

(iii) We derive upper bounds on the performance of our nonparametric and parametric pricing policies (see Propositions 1 and 3). In all cases, the proposed policies achieve a regret that is “not far” from the lower bound described above. In the parametric setting when only one parameter is unknown, we prove that essentially no admissible pricing policy can achieve a smaller regret than our proposed method (see Proposition 5).

(iv) Building on ideas from stochastic approximations, we indicate how one can develop more refined sequential policies, and illustrate this in the nonparametric setting (see Algorithm 4 in Online Appendix A).

Returning to the questions raised earlier in this section, our results shed light on the following issues. First, despite having only limited (or almost no) prior information, it is possible to construct joint learning and pricing policies that generate revenues that are “close” to the best achievable performance with full information. Our results highlight an interesting observation. In the full-information setting, Gallego and van Ryzin (1994) prove that fixed-price heuristics lead to near-optimal revenues; hence, the value of dynamic price changes is (at least asymptotically) limited. In a setting with incomplete information, price changes play a much more pivotal role because they are relied upon to resolve uncertainty with regard to the demand function.

The regret bounds described above rigorously quantify the economic value of a priori information on the demand model. Alternatively, the lost revenues can be viewed as quantifying the “price” paid due to model uncertainty. Finally, our work highlights an important issue related to model misspecification risk. In particular, if an algorithm is designed under parametric assumptions, it is prone to such risk because the true demand function may not (and in many cases will not) belong to the assumed parametric family. Our regret bounds provide a means for quantifying the “price” that one pays for eliminating this risk via nonparametric approaches; see also the numerical illustration in §6.

The Remainder of This Paper. The next section reviews related literature. Section 3 introduces the model and formulates the problem. Section 4 studies the nonparametric setting, and §5 focuses on cases where the demand function possesses a parametric structure. Section 6 presents numerical results and discusses some qualitative insights. Section 7 provides proofs of Propositions 1–3. This paper

has an online supplement with three parts, which can be found at <http://or.pubs.informs.org/>. Appendix A formulates adaptive versions of the nonparametric algorithm and illustrates their performance. Appendix B contains the proof of Proposition 5 and Appendix C contains auxiliary lemmas and their proofs.

2. Related Literature

Parametric Approaches. The majority of revenue management studies that address demand function uncertainty do so by assuming that one or more parameters characterizing this function are unknown. The typical approach here follows a dynamic programming formulation with Bayesian updating, where a prior on the distribution of the unknown parameters is initially postulated. Recent examples include Aviv and Pazgal (2005), Araman and Caldentey (2005), and Farias and Van Roy (2006), all of which assume that a single parameter is unknown. (See also Lobo and Boyd 2003 and Carvalho and Puterman 2005.) Scarf (1959) was one of the first papers to use this Bayesian formulation, although in the context of inventory management.

Although the Bayesian approach provides for an attractive stylized analysis of the joint learning and pricing problem, it suffers from significant shortcomings. Most notably, the objective of the dynamic optimization problem involves an expectation that is taken relative to a prior distribution over the unknown parameters. Hence, any notion of optimality associated with a Bayesian-based policy is with respect to that prior. Moreover, the specification of this prior distribution is typically constrained to so-called conjugate families and is not driven by “real” prior information; the hindering element here is the computation of the posterior via Bayes rule. The above factors introduce significant restrictions on the models that are amenable to analysis via Bayesian dynamic programming. Bertsimas and Perakis (2006) have considered an alternative to this formulation using a least squares approach in the context of a linear demand model. Our work in the parametric setting is based on a frequentist approach and hence is applicable to a wide class of parametric models; as such, it offers an alternative to Bayesian approaches that circumvents some of their deficiencies.

Nonparametric Approaches. The main difficulty facing nonparametric approaches is loss of tractability. Most work here has been pursued in relatively simple static settings that do not allow for learning of the demand function; see, e.g., Ball and Queyranne (2009) and Eren and Maglaras (2007) for a competitive ratio formulation, and Perakis and Roels (2009) for a minimax regret formulation. These studies focus almost exclusively on structural insights. The recent paper by Rusmevichientong et al. (2006) develops a nonparametric approach to a multiproduct static pricing problem, based on historical data. The formulation does not incorporate inventory constraints or

finite sales horizon considerations. In the context of optimizing seat allocation policies for a single-flight multiclass problem, van Ryzin and McGill (2000) show that one can use stochastic approximation methods to reach near-optimal capacity protection levels in the long run. The absence of parametric assumptions in this case is with respect to the distributions of customers’ requests for each class. (See also Huh and Rusmevichientong 2008 for a related study.)

Perhaps the paper most closely related to our current work is that of Lim and Shanthikumar (2007), who formulate a robust counterpart to the single-product revenue management problem of Gallego and van Ryzin (1994). In that paper the uncertainty arises at the level of the point process distribution characterizing realized demand, and the authors use a max-min formulation where nature is adversarial at *every point in time*. This type of conservative setting effectively precludes any real-time learning and, moreover, does not lend itself to prescriptive solutions.

Related Work in Other Disciplines. The general problem of dynamic optimization with limited or no information about a response function has also attracted attention in other fields. In economics, a line of work that traces back to Hannan (1957) studies settings where the decision maker faces an oblivious opponent. The objective is to minimize the difference between the rewards accumulated by a given policy, and the rewards accumulated by the best-possible single action had the decision maker known in advance the actions of the adversary; see Foster and Vohra (1999) for a review of this line of work and its relation to developments in other fields.

A classical formulation of sequential optimization under uncertainty that captures the essence of the exploration-exploitation trade-off is the multiarmed bandit paradigm that dates back to the work of Robbins (1952). This was originally introduced as a model of clinical trials in the statistics literature, but has since been used in many other settings; see, e.g., Lai and Robbins (1985) and references therein. Related studies in the computer science literature include Auer et al. (2002), who study an adversarial version of a multiarmed bandit problem, and Kleinberg and Leighton (2003), who provide an analysis of an online posted-price auction using these tools. (See Cesa-Bianchi and Lugosi 2006 for a recent and comprehensive survey.) Our work shares an important common theme with the streams of literature surveyed above, insofar as it, too, highlights exploration-exploitation trade-offs. On the other hand, our work differs significantly from antecedent literature along three dimensions that are characteristic of our dynamic pricing problem: we deal with a *constrained* dynamic optimization problem (the constraint arising from the initial inventory level); the action space of the decision maker, namely, the feasible price set, is uncountable; and the action space of the adversary (nature), namely, the class of admissible demand functions, is also uncountable.

3. Problem Formulation

Model Primitives and Basic Assumptions. We consider a revenue management problem in which a monopolist sells a single product. The selling horizon is denoted by $T > 0$, and after this time sales are discontinued and there is no salvage value for the remaining unsold products. Demand for the product at any time $t \in [0, T]$ is given by a Poisson process with intensity λ_t that measures the instantaneous demand rate (in units such as number of products requested per hour, say): letting $\lambda: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ denote the *demand function*, then if the price at time t is $p(t)$, the instantaneous demand rate at time t is given by $\lambda_t = \lambda(p(t))$, and realized demand is a controlled Poisson process with this intensity.

We assume that the set of feasible prices is $[\underline{p}, \bar{p}] \cup p_\infty$, where $0 < \underline{p} < \bar{p} < \infty$ and $p_\infty > 0$ is a price that “turns off” demand (and revenue rate), i.e., $\lambda(p_\infty) = 0$. (The case $p_\infty = \infty$ can be treated by assuming $r(\lambda(p_\infty)) := \lim_{p \rightarrow p_\infty} r(\lambda(p)) = 0$, which yields $\lim_{p \rightarrow p_\infty} \lambda(p) = 0$.) With regard to the demand function, we assume that $\lambda(\cdot)$ is nonincreasing in the price p , has an inverse denoted by $\gamma(\cdot)$, and the revenue rate $r(\lambda) := \lambda\gamma(\lambda)$ is concave. These assumptions are quite standard in the revenue management literature, resulting in the term *regular* affixed to demand functions satisfying these conditions; see, e.g., Talluri and van Ryzin (2005, §7).

Let $(p(t): 0 \leq t \leq T)$ denote the *price process*, which is assumed to have sample paths that are right continuous, with left limits taking values in $[\underline{p}, \bar{p}] \cup p_\infty$. Let $N(\cdot)$ be a unit-rate Poisson process. The cumulative demand for the product up until time t is then given by $D(t) := N(\int_0^t \lambda(p(s)) ds)$. We say that $(p(t): 0 \leq t \leq T)$ is *nonanticipating* if the value of $p(t)$ at each time $t \in [0, T]$ is only allowed to depend on past prices $\{p(s): s \in [0, t]\}$ and demand values $\{(D(s)): s \in [0, t]\}$. (More formally, the price process is adapted to the filtration generated by the past values of the demand and price processes.)

Information Structure and the Economic Optimization Problem. We assume that the decision maker *does not know* the true demand function λ , but is able to continuously observe realized demand at all time instants, starting at time 0 and up until the end of the selling horizon T . The only information available regarding λ is that it belongs to a class of admissible demand functions, \mathcal{L} ; in §4, \mathcal{L} will be taken to be a *nonparametric* class of functions, and in §5 it will be restricted to a *parametric* class. Thus, the makeup of the class \mathcal{L} summarizes prior information on the demand model.

We shall use π to denote a *pricing policy*, which, roughly speaking, maps the above information structure to a nonanticipating price process $(p(t): 0 \leq t \leq T)$. With some abuse of terminology, we will use the term “policy” to refer to the price process itself and the algorithm that generates it, interchangeably. Put

$$N^\pi(t) := N\left(\int_0^t \lambda(p(s)) ds\right) \quad \text{for } 0 \leq t \leq T, \quad (1)$$

where $N^\pi(t)$ denotes the cumulative demand up to time t under the policy π .

Let $x > 0$ denote the inventory level (number of products) at the start of the selling season. A pricing policy π is said to be *admissible* if the induced price process satisfies

$$\int_0^T dN^\pi(s) \leq x \quad \text{a.s.}, \quad (2)$$

$$p(s) \in [\underline{p}, \bar{p}] \cup p_\infty, \quad 0 \leq s \leq T. \quad (3)$$

It is important to note that although the decision maker does not know the demand function, knowledge that $\lambda(p_\infty) = 0$ guarantees that the constraint (2) can be met. Let \mathcal{P} denote the set of admissible pricing policies.

The dynamic optimization problem faced by the decision maker under the information structure described above is: choose $\pi \in \mathcal{P}$ to maximize the total expected revenues

$$J^\pi(x, T; \lambda) := \mathbb{E}\left[\int_0^T p(s) dN^\pi(s)\right]. \quad (4)$$

The dependence on λ in the left-hand side is indicative of the fact that the expectation on the right-hand side is taken with respect to the *true* demand distribution. Because the decision maker cannot compute the expectation in (4) without knowing the underlying demand function, the above optimization problem does not seem to be well posed. In a sense, one can view the solution of (4) as being made possible only with the aid of an “oracle” that can compute the quantity $J^\pi(x, T; \lambda)$ for any given policy. We will now redefine the decision maker’s objective in a more suitable manner, using the notion of a full-information benchmark.

A Full-Information Benchmark. Let us first explain how the analysis of the dynamic optimization problem described in (4) proceeds when one removes two significant obstacles: lack of knowledge of the demand function λ prior to the start of the selling season, and stochastic variability in realized demand. In particular, consider the following full-information deterministic optimization problem, in which the function λ is assumed to be known at time $t = 0$:

$$\begin{aligned} & \sup \int_0^T r(\lambda(p(s))) ds, \\ & \text{s.t. } \int_0^T \lambda(p(s)) ds \leq x, \\ & \quad p(s) \in [\underline{p}, \bar{p}] \cup p_\infty \quad \text{for all } s \in [0, T]. \end{aligned} \quad (5)$$

This problem is obtained from (4), and the admissibility conditions (2) and (3), by replacing the random process characterizing customer purchase requests by its mean rate. For example, if one focuses on the objective (4), then the deterministic objective in (5) is obtained by substituting $\lambda(p(s))ds$ for $dN^\pi(s)$ because $r(\lambda(p(s))) = p(s)\lambda(p(s))$. The same parallel can be drawn between the first constraint of the deterministic problem and (2). Consequently, it is

reasonable to refer to (5) as a *full-information deterministic relaxation* of the original dynamic pricing problem (4).

Let us denote the value of (5) as $J^D(x, T | \lambda)$, where D is mnemonic for deterministic, and the choice of notation with respect to λ reflects the fact that the optimization problem is solved conditioned on knowing the true underlying demand function. The value of the full-information deterministic relaxation provides, as one would anticipate, an upper bound on expected revenues generated by any pricing policy $\pi \in \mathcal{P}$, that is, $J^\pi(x, T; \lambda) \leq J^D(x, T | \lambda)$ for all $\lambda \in \mathcal{L}$. This rather intuitive observation is formalized in Lemma 1 in Online Appendix C, which essentially generalizes Gallego and van Ryzin (1994, Proposition 2).

The Minimax Regret Objective. As indicated above, for any demand function $\lambda \in \mathcal{L}$, we have that $J^\pi(x, T; \lambda) \leq J^D(x, T | \lambda)$ for all admissible policies $\pi \in \mathcal{P}$. With this in mind, we define the *regret* $\mathcal{R}^\pi(x, T; \lambda)$ for any given function $\lambda \in \mathcal{L}$ and policy $\pi \in \mathcal{P}$ to be

$$\mathcal{R}^\pi(x, T; \lambda) = 1 - \frac{J^\pi(x, T; \lambda)}{J^D(x, T | \lambda)}. \quad (6)$$

The regret measures the percentage loss in performance of any policy π in relation to the benchmark $J^D(x, T | \lambda)$. By definition, the value of the regret always lies in the interval $[0, 1]$, and the smaller the regret, the better the performance of a policy π ; in the extreme case when the regret is zero, then the policy π is guaranteed to extract the maximum full-information revenues. Because the decision maker does not know which demand function s/he will face in the class \mathcal{L} , it is attractive to design pricing policies that perform well irrespective of the actual underlying demand function. In particular, if the decision maker uses a policy $\pi \in \mathcal{P}$, and nature then “picks” the worst possible demand function for that policy, then the resulting regret would be

$$\sup_{\lambda \in \mathcal{L}} \mathcal{R}^\pi(x, T; \lambda). \quad (7)$$

In this game-theoretic setting it is now possible to restate the decision maker’s objective, initially given in (4), as follows: pick $\pi \in \mathcal{P}$ to minimize (7). The advantage of this formulation is that the decision maker’s problem is now well posed: for any $\pi \in \mathcal{P}$ and fixed $\lambda \in \mathcal{L}$, it is possible, at least in theory, to compute the numerator on the right-hand side in (6), and hence (7). Roughly speaking, one can attach a worst case $\lambda \in \mathcal{L}$ to each policy $\pi \in \mathcal{P}$, and subsequently one can try to optimize this by searching for the policy with the best worst-case performance. In other words, we are interested in characterizing the *minimax regret*

$$\inf_{\pi \in \mathcal{P}} \sup_{\lambda \in \mathcal{L}} \mathcal{R}^\pi(x, T; \lambda). \quad (8)$$

This quantity has an obvious physical interpretation: it measures the monetary value (in normalized currency units) of knowing the demand function a priori. The issue, of course,

is that, barring exceedingly simple cases, it is not possible to compute the minimax regret. Our objective in what follows will be to characterize this quantity by deriving suitable bounds on (8) in cases where the class \mathcal{L} is nonparametric or is restricted to a suitable parametric class of demand functions.

4. Main Results: The Nonparametric Case

4.1. A Nonparametric Pricing Algorithm

We introduce below a learning and pricing policy defined through two tuning parameters (κ, τ) : κ is a positive integer and $\tau \in (0, T]$. The general structure, which is summarized for convenience in algorithmic form, is divided into two main stages. A “learning” phase (exploration) of length τ is used first, in which κ prices are tested. Then, a “pricing” phase (exploitation) fixes a “good” price based on demand observations in the first phase. The intuition underlying the method is discussed immediately following the description of the method.

ALGORITHM 1. $\pi(\tau, \kappa)$

Step 1. Initialization:

(a) Set the learning interval to be $[0, \tau]$, and the number of prices to experiment with to be κ . Put $\Delta = \tau/\kappa$.

(b) Divide $[\underline{p}, \bar{p}]$ into κ equally spaced intervals and let $\{p_i, i = 1, \dots, \kappa\}$ be the left endpoints of these intervals.

Step 2. Learning/experimentation:

(a) On the interval $[0, \tau]$ apply p_i from $t_{i-1} = (i-1)\Delta$ to $t_i = i\Delta$, $i = 1, 2, \dots, \kappa$, as long as inventory is positive. If no more units are in stock, apply p_∞ up until time T and STOP.

(b) Compute

$$\hat{d}(p_i) = \frac{\text{total demand over } [t_{i-1}, t_i]}{\Delta}, \quad i = 1, \dots, \kappa.$$

Step 3. Optimization:

$$\begin{aligned} \text{Compute } \hat{p}^u &= \arg \max_{1 \leq i \leq \kappa} \{p_i \hat{d}(p_i)\}, \\ \hat{p}^c &= \arg \min_{1 \leq i \leq \kappa} |\hat{d}(p_i) - x/T|, \end{aligned} \quad (9)$$

$$\text{and set } \hat{p} = \max\{\hat{p}^c, \hat{p}^u\}. \quad (10)$$

Step 4. Pricing:

On the interval $(\tau, T]$ apply \hat{p} as long as inventory is positive, then apply p_∞ for the remaining time.

Intuition and Key Underlying Ideas. At first, a nonparametric empirical estimate of the demand function is obtained based on a learning phase of length τ described in Steps 1 and 2. The intuition underlying Steps 3 and 4 is based on the analysis of the deterministic relaxation (5) whose solution (see Lemma 1 in Online Appendix C)

is given by $p(s) = p^D := \max\{p^u, p^c\}$ for $s \in [0, T']$ and $p(s) = p_\infty$ for $s > T'$, where

$$p^u = \arg \max_{p \in [\underline{p}, \bar{p}]} \{r(\lambda(p))\}, \quad p^c = \arg \min_{p \in [\underline{p}, \bar{p}]} |\lambda(p) - x/T|,$$

and $T' = \min\{T, x/\lambda(p^D)\}$. Here, the superscripts u and c stand for *unconstrained* and *constrained*, respectively, in reference to whether or not the inventory constraint in (5) is binding. In particular, the deterministic problem (5) can be solved by restricting attention to the two prices p^u and p^c . Algorithm 1 hinges on this observation.

In Step 3, the objective is to obtain an accurate estimate of p^D based on the observations during the exploration phase, while at the same time keeping τ small to limit the revenue loss over this learning phase. The algorithm then applies this price on $(\tau, T]$. With the exception of the short initial phase $[0, \tau]$, the expected revenues (in the real system) will be close to those achieved by p^D over $[0, T]$. The analysis in Gallego and van Ryzin (1994) establishes that those revenues would be close to $J^D(x, T | \lambda)$, and as a result, the regret $\mathcal{R}^\pi(x, T; \lambda)$ should be small.

To estimate p^D , the algorithm dedicates an initial portion $[0, \tau]$ of the total selling interval $[0, T]$ to an exploration of the price domain. On this initial interval, the algorithm experiments with κ prices where each is kept fixed for τ/κ units of time. This structure leads to three main sources of error in the search for p^D . First, during the learning phase one incurs an *exploration bias* because the prices being tested there are not close to p^D (or close to the optimal fixed price for that matter). This incurs losses of order τ . Second, experimenting with only a finite number of prices κ in the search for p^D results in a *deterministic error* of order $1/\kappa$ in Step 3. Finally, only “noisy” demand observations are available at each of the κ price points, and the longer a price is held fixed, the more accurate the estimate of the mean demand rate at that price. This introduces a *stochastic error* of order $(\tau/\kappa)^{-1/2}$ stemming from the nature of the Poisson process. The crux of the matter is to balance these three error sources by a suitable choice of the tuning parameters τ and κ .

4.2. Model Uncertainty: The Class of Demand Functions

The nonparametric class of functions that we consider consists of regular demand functions (satisfying the standard conditions laid out in §3), which, in addition, satisfy the following:

ASSUMPTION 1. For some finite positive constants M , \underline{K} , \bar{K} , m , with $\underline{K} \leq \bar{K}$:

- (i) *Boundedness*: $|\lambda(p)| \leq M$ for all $p \in [\underline{p}, \bar{p}]$.
- (ii) *Lipschitz continuity*: $|\lambda(p) - \lambda(p')| \leq \bar{K}|p - p'|$ for all $p, p' \in [\underline{p}, \bar{p}]$ and $|\gamma(l) - \gamma(l')| \leq \underline{K}^{-1}|l - l'|$ for all $l, l' \in [\lambda(\bar{p}), \lambda(\underline{p})]$.
- (iii) *Minimum revenue rate*: $\max\{p\lambda(p) : p \in [\underline{p}, \bar{p}]\} \geq m$.

Let $\mathcal{L} := \mathcal{L}(M, \underline{K}, \bar{K}, m)$ denote this class of demand functions. Assumptions 1(i) and 1(ii) are quite benign, only requiring minimal smoothness of the demand function; (ii) also ensures that when $\lambda(\cdot)$ is positive, it does not have “flat regions.” Assumption 1(iii) states that a minimal revenue rate exists, hence avoiding trivialities.

Note that Assumptions (i)–(iii) above hold for many models of the demand function used in the revenue management and economics literature (e.g., linear, exponential, and isoelastic/Pareto with parameters lying in a compact set; cf. Talluri and van Ryzin 2005, §7 for further examples).

4.3. Performance Analysis

Because minimax regret is hardly a tractable quantity, we introduce in this section an asymptotic regime characterized by a high volume of sales, which will be used to analyze the performance of Algorithm 1. We consider a regime in which both the size of the initial inventory as well as potential demand grow proportionally large. In particular, for a market of size n , where n is a positive integer, the initial inventory and the demand function are now assumed to be given by

$$x_n = nx \quad \text{and} \quad \lambda_n(\cdot) = n\lambda(\cdot). \quad (11)$$

Thus, the index n determines the order of magnitude of both inventory and rate of demand. We will denote by \mathcal{P}_n the set of admissible policies for a market of scale n , and the expected revenues under a policy $\pi_n \in \mathcal{P}_n$ will be denoted $J_n^\pi(x, T; \lambda)$. With some abuse of notation, we will occasionally use π to denote a sequence of policies $\{\pi_n, n = 1, 2, \dots\}$ as well as any element of that sequence, omitting the subscript n to avoid cluttering the notation. For each $n = 1, 2, \dots$, we denote by $J_n^D(x, T | \lambda)$ the value of the deterministic relaxation given in (5) with the scaling given in (11); it is straightforward to verify that $J_n^D(x, T | \lambda) = nJ^D(x, T | \lambda)$. Finally, let the regret be denoted as $\mathcal{R}_n^\pi(x, T; \lambda) := 1 - J_n^\pi(x, T; \lambda)/J_n^D(x, T | \lambda)$. The following definition characterizes admissible policies that have “good” asymptotic properties.

DEFINITION 1 (ASYMPTOTIC OPTIMALITY). A sequence of admissible policies $\pi_n \in \mathcal{P}_n$ is said to be asymptotically optimal if

$$\sup_{\lambda \in \mathcal{L}} \mathcal{R}_n^\pi(x, T; \lambda) \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (12)$$

In other words, asymptotically optimal policies achieve the full-information upper bound on revenues as $n \rightarrow \infty$, *uniformly* over the class of admissible demand functions.

For the purpose of asymptotic analysis, we use the following notation: for real-valued positive sequences $\{a_n\}$ and $\{b_n\}$, we write $a_n = O(b_n)$ if a_n/b_n is bounded from above by a constant, and if a_n/b_n is also bounded from below, then we write $a_n \asymp b_n$. We now analyze the performance of policies associated with Algorithm 1 in the asymptotic regime described above.

PROPOSITION 1. *Let Assumption 1 hold. Set $\tau_n \asymp n^{-1/4}$, $\kappa_n \asymp n^{1/4}$ and let $\pi_n := \pi(\tau_n, \kappa_n)$ be given by Algorithm 1. Then, the sequence $\{\pi_n\}$ is asymptotically optimal, and for all $n \geq 1$,*

$$\sup_{\lambda \in \mathcal{L}} \mathcal{R}_n^\pi(x, T; \lambda) \leq \frac{C(\log n)^{1/2}}{n^{1/4}} \quad (13)$$

for some finite positive constant C .

The constant C above depends only on the parameters characterizing the class \mathcal{L} , the initial inventory x , and the time horizon T . The exact dependence is somewhat complex and is omitted; however, we note that it is fully consistent with basic intuition: as one expands the class \mathcal{L} by suitably increasing or decreasing the value of the parameters in Assumption 1, the magnitude of C grows, and vice versa. (This can also be inferred by carefully inspecting the proof of the proposition.)

We next present a lower bound on the minimax regret which establishes a fundamental limit on the performance of any admissible pricing policy. Roughly speaking, the main idea behind this result is to construct a worst-case demand function such that the regret is large for any policy.

PROPOSITION 2. *Let Assumption 1 hold with M, \bar{K} satisfying $M \geq \max\{2\bar{K}\bar{p}, \bar{K}\bar{p} + x/T\}$. Then, there exists a finite positive constant C' such that for any sequence of admissible policies $\{\pi_n\}$ and for all $n \geq 1$,*

$$\sup_{\lambda \in \mathcal{L}} \mathcal{R}_n^\pi(x, T; \lambda) \geq \frac{C'}{n^{1/2}}. \quad (14)$$

Combining Propositions 1 and 2, one can characterize the magnitude of the minimax regret as follows:

$$\frac{C'}{n^{1/2}} \leq \inf_{\pi \in \mathcal{P}_n} \sup_{\lambda \in \mathcal{L}} \mathcal{R}_n^\pi(x, T; \lambda) \leq \frac{C(\log n)^{1/2}}{n^{1/4}}, \quad (15)$$

and hence the performance of Algorithm 1 is “not far” from being minimax optimal (i.e., achieving the lower bound). A question that remains open is whether one can close this gap by further refining the algorithm. We revisit this point in §5 and Online Appendix A.

5. Main Results: The Parametric Case

In this section, we assume that the demand function is known to have a parametric form. Our goal is to develop pricing policies that exploit this information and that will work well for a large class of admissible parametric demand functions. One of the main questions of interest here is whether these policies achieve a smaller regret relative to the nonparametric case, and if there exist policies whose performance cannot be improved upon.

Preliminaries. Let k denote a positive integer and $\mathcal{L}(\Theta) = \{\lambda(\cdot; \theta) : \theta \in \Theta\}$ be a parametric family of demand

functions, where $\Theta \subseteq \mathbb{R}^k$ is assumed to be a convex compact set, $\theta \in \Theta$ is a parameter vector, and $\lambda: \mathbb{R}_+ \times \Theta \rightarrow \mathbb{R}_+$. We consider all parametric families that are subsets of the class of admissible regular demand functions defined in §4, i.e., $\mathcal{L}(\Theta) \subset \mathcal{L} := \mathcal{L}(M, \underline{K}, \bar{K}, m)$.

In the current setting, we will assume that the decision maker knows that $\lambda \in \mathcal{L}(\Theta)$, i.e., s/he knows the parametric structure of the demand function, but *does not know* the value of the parameter vector θ . We continue to denote by \mathcal{P} the set of admissible pricing policies, i.e., policies that satisfy (2) and (3). For any policy $\pi \in \mathcal{P}$, let $J^\pi(x, T; \theta)$ denote the expected revenues under π , and let $J^D(x, T | \theta)$ denote the value of the deterministic relaxation (5) when the value of the unknown parameter vector is revealed to the decision maker prior to the start of the selling season. Let $\mathcal{R}^\pi(x, T; \theta)$ denote the regret under a policy π , namely,

$$\mathcal{R}^\pi(x, T; \theta) = 1 - \frac{J^\pi(x, T; \theta)}{J^D(x, T | \theta)}. \quad (16)$$

5.1. The Proposed Method

We consider a simple modification of the approach taken in Algorithm 1 that now exploits the assumed parametric structure of the demand function.

ALGORITHM 2. $\pi(\tau)$

Step 1. Initialization:

(a) Set the learning interval to be $[0, \tau]$. Put $\Delta = \tau/k$, where k is the dimension of the parameter space Θ .

(b) Choose a set of k prices, $\{p_i, i = 1, \dots, k\}$.

Step 2. Learning/experimentation:

(a) On the interval $[0, \tau]$ apply p_i from $t_{i-1} = (i-1)\Delta$ to $t_i = i\Delta$, $i = 1, \dots, k$, as long as the inventory is positive. If no more units are in stock, apply p_∞ up until time T and STOP.

(b) Compute

$$\hat{d}_i = \frac{\text{total demand over } [t_{i-1}, t_i]}{\Delta}, \quad i = 1, \dots, k.$$

(c) Let $\hat{\theta}$ be a solution of $\{\lambda(p_i; \theta) = \hat{d}_i, i = 1, \dots, k\}$.

Step 3. Optimization:

$$p^u(\hat{\theta}) = \arg \max\{p\lambda(p; \hat{\theta}) : p \in [\underline{p}, \bar{p}]\},$$

$$p^c(\hat{\theta}) = \arg \min\{|\lambda(p; \hat{\theta}) - x/T| : p \in [\underline{p}, \bar{p}]\},$$

$$\hat{p} = \max\{p^u(\hat{\theta}), p^c(\hat{\theta})\}.$$

Step 4. Pricing:

On the interval $(\tau, T]$ apply \hat{p} as long as inventory is positive, then apply p_∞ for the remaining time.

Note that Step 1(b) requires one to define the prices $\{p_1, \dots, p_k\}$ and Step 2(c) implicitly assumes that the system of equations admits a solution. (We define the prices and state this assumption more formally when analyzing

the performance of Algorithm 2.) The intuition behind this algorithm is similar to the one that underlies the construction of Algorithm 1, the only difference being that the parametric structure allows one to infer accurate information about the demand function using only a small number of test prices (k). In particular, recalling the discussion following Algorithm 1, a key difference is that the deterministic error source associated with price granularity does not affect the performance of Algorithm 2.

5.2. The Parametric Class of Demand Functions

First, note that the inclusion $\mathcal{L}(\Theta) \subset \mathcal{L}$ implies that $\lambda(\cdot; \theta)$ has an inverse for all $\theta \in \Theta$; this inverse will be denoted $\gamma(\cdot; \theta)$. (Note also that conditions (i)–(iii) in Assumption 1 hold.) For any parameter vector $\theta \in \Theta$, we denote the revenue function by $r(l; \theta) := \gamma(l; \theta)l$. Let $\{P_\theta: \theta \in \Theta\}$ be the family of demand distributions corresponding to Poisson processes with controlled intensities $\lambda(\cdot; \theta), \theta \in \Theta$. We denote by $f_{\lambda(p; \theta)}(\cdot)$ the probability mass function of a Poisson random variable with intensity $\lambda(p; \theta)$.

The following technical conditions articulate standard regularity assumptions in the context of maximum-likelihood estimation (cf. Borovkov 1998) and are used to define the admissible class of parametric demand functions.

ASSUMPTION 2. (i) *There exists a vector of distinct prices $\bar{p} = (p_1, \dots, p_k) \in [\underline{p}, \bar{p}]$ such that:*

- (a) *For some $l_0 > 0$, $\min_{1 \leq i \leq k} \inf_{\theta \in \Theta} \lambda(p_i; \theta) > l_0$.*
- (b) *For any vector $\vec{d} = (d_1, \dots, d_k)$, the system of equations $\{\lambda(p_i; \theta) = d_i, i = 1, \dots, k\}$ has a unique solution in θ . Let $g(\vec{p}, \vec{d})$ denote this solution. We assume in addition that $g(\vec{p}, \cdot)$ is Lipschitz continuous with constant $\alpha > 0$.*
- (c) *For $i = 1, \dots, k$, $\sqrt{\lambda(p_i; \theta)}$ is differentiable on Θ .*

(ii) *For some $\bar{K}_2 > 0$, $|\lambda(p; \theta) - \lambda(p; \theta')| \leq \bar{K}_2 \|\theta - \theta'\|_\infty$ for all $p \in [\underline{p}, \bar{p}]$ and $\theta, \theta' \in \Theta$.*

Condition (i) ensures that the parametric model is identifiable based on a sufficient set of observations. Condition (ii) is a mild regularity assumption on the parametric class, controlling for changes in the demand curve as parameters vary. We provide below an example of a parametric class that satisfies Assumption 2.

EXAMPLE 1 (LINEAR DEMAND FUNCTION). Let $\lambda(p; \theta) = \theta_1 - \theta_2 p$, and set $[\underline{p}, \bar{p}] = [1, 2]$. Let $\mathcal{L}(\Theta) = \{\lambda(\cdot; \theta): \theta \in [10, 20] \times [1, 4]\}$. If one sets $(p_1, p_2) = (1, 2)$, it is straightforward to verify conditions (i)(a), (i)(c), and that (ii) is satisfied with $\bar{K}_2 = 1$. It is also easy to see that the unique solution associated with condition (i)(b) is given by

$$g(p_1, p_2, d_1, d_2) = \left(d_1 + \frac{p_1}{p_2 - p_1} [d_1 - d_2], \frac{1}{p_2 - p_1} [d_1 - d_2] \right),$$

and $g(p_1, p_2, d_1, d_2)$ is clearly Lipschitz continuous with respect to (d_1, d_2) . \square

5.3. Performance Analysis

Suppose that the initial inventory level and the demand function are scaled according to (11), and denote the regret by $\mathcal{R}_n^\pi(x, T; \theta) = 1 - J_n^\pi(x, T; \theta) / J_n^D(x, T | \theta)$.

PROPOSITION 3. *Let Assumptions 1 and 2 hold and let $\{p_1, \dots, p_k\}$ be as in Assumption 2(i). Set $\tau_n \asymp n^{-1/3}$ and let $\pi_n := \pi(\tau_n)$ be defined by Algorithm 2. Then, the sequence of policies $\{\pi_n\}$ is asymptotically optimal and satisfies*

$$\sup_{\theta \in \Theta} \mathcal{R}_n^\pi(x, T; \theta) \leq \frac{C(\log n)^{1/2}}{n^{1/3}} \tag{17}$$

for all $n \geq 1$ and some finite positive constant C .

Contrasting the above with Proposition 1, we observe that the parametric structure of the demand function translates into improved performance bounds for the learning and pricing algorithm that is designed with this knowledge in mind. In particular, the regret is now of order $\mathcal{R}_n = O((\log n)^{1/2} n^{-1/3})$ as opposed to $\mathcal{R}_n = O((\log n)^{1/2} n^{-1/4})$ in the nonparametric case. In addition, note that the upper bound above holds for *all* admissible parametric families. The improvement in terms of generated revenues, as quantified by the smaller magnitude of the regret, spells out the advantages of using a parametric approach versus a nonparametric one. The downside—namely, model misspecification and its consequences—is illustrated and discussed in §6.

At an intuitive level, classical estimation theory tells us that parameter uncertainty cannot be resolved faster than rate $n^{-1/2}$ with n observations. This suggests that for the asymptotic regime we consider in this paper, no admissible pricing policy would be able to achieve a convergence rate faster than $n^{-1/2}$. This intuition is made rigorous in the next proposition.

PROPOSITION 4. *Let Assumption 1 hold with M, \bar{K} satisfying $M \geq \max\{2\bar{K}\bar{p}, \bar{K}\bar{p} + x/T\}$. Then, there exists a parametric family $\mathcal{L}(\Theta) \subset \mathcal{L}(M, \underline{K}, \bar{K}, m)$ satisfying Assumption 2 such that for some positive constant C' and for all admissible policies $\{\pi_n\}$,*

$$\sup_{\theta \in \Theta} \mathcal{R}_n^\pi(x, T; \theta) \geq \frac{C'}{n^{1/2}} \tag{18}$$

for all $n \geq 1$.

This result follows in a relatively straightforward manner from Proposition 2.

5.4. An Optimal Algorithm When a Single Parameter Is Unknown

Given the lower bound in (18), the natural question that arises is whether it is possible to close the remaining gap with respect to the upper bound in Proposition 3. We study this question in the context of a single unknown parameter, i.e., $k = 1$, and hence $\Theta \subseteq \mathbb{R}$. Consider the following ℓ -step policy $\pi(\ell, \Delta^{(1)}, \dots, \Delta^{(\ell)})$.

ALGORITHM 3. $\pi(\ell, \Delta^{(1)}, \dots, \Delta^{(\ell)})$

Step 1. Initialization:

(a) Set the number of steps to be ℓ and define $\Delta^{(i)}$, $i = 1, \dots, \ell$ so that $\Delta^{(1)} + \dots + \Delta^{(\ell)} = T$.

(b) Choose a price $\hat{p}_1 \in [p, \bar{p}]$ as in Assumption 2(i).

Step 2. Learning/Optimization/Pricing:

Set $t_1 = 0$.

For $i = 1, \dots, \ell$,

(a) Learning/Pricing:

(i) Apply \hat{p}_i on the interval $[t_i, t_i + \Delta^{(i)})$ as long as the inventory is positive. If no more units are in stock, apply p_∞ up until time T and STOP.

(ii) Compute

$$\hat{d}_i = \frac{\text{total demand over } [t_i, t_i + \Delta^{(i)})}{\Delta^{(i)}}.$$

(iii) Set $t_{i+1} = t_i + \Delta^{(i)}$.

(iv) Let $\hat{\theta}^i$ be the unique solution of $\lambda(p_i, \theta) = \hat{d}_i$.

(b) Optimization:

$$p^u(\hat{\theta}^i) = \arg \max \{p \lambda(p; \hat{\theta}^i) : p \in [p, \bar{p}]\},$$

$$p^c(\hat{\theta}^i) = \arg \min \{|\lambda(p; \hat{\theta}^i) - x/t| : p \in [p, \bar{p}]\},$$

$$\hat{p}_{i+1} = \max \{p^u(\hat{\theta}^i), p^c(\hat{\theta}^i)\}.$$

End For

The intuition underlying Algorithm 3 is as follows. When a single parameter is unknown, one can infer information about the parameter from observations of demand at a single price. Given this, the idea of Algorithm 3 is to price close to p^D after the first stage, but to *continue* learning. In particular, the estimate of p^D is improved from stage to stage using the demand observations from the previous stage. As a result, losses are mitigated by two effects: (i) after stage 1, the price is always close to p^D ; and (ii) the estimate of p^D becomes more precise. For what follows, we slightly strengthen Assumption 2(i).

ASSUMPTION 3. $\inf_{p \in [p, \bar{p}]} \inf_{\theta \in \Theta} \lambda(p; \theta) > l_0$ and for any price $p \in [p, \bar{p}]$ and any $d \geq 0$, the equation $\lambda(p; \cdot) = d$ has a unique solution. If $g(p, d)$ denotes this solution, then $g(p, \cdot)$ is Lipschitz continuous with constant $\alpha > 0$.

We now analyze the performance of policies associated with Algorithm 3. In particular, suppose that the initial inventory level and the demand function are scaled according to (11), and define the sequence of tuning parameters $\{\ell_n, \Delta_n^{(1)}, \dots, \Delta_n^{(\ell_n)}\}$ as follows:

$$\ell_n = (\log 2)^{-1} \log \log n, \quad (19)$$

$$\Delta_n^{(m)} = \beta_n n^{(a_m/a_m)^{-1}}, \quad m = 1, \dots, \ell_n, \quad (20)$$

where $a_m = 2^{m-1}/(2^m - 1)$ for $m \geq 1$, and $\beta_n > 0$ is a normalizing constant chosen so $\Delta_n^{(1)} + \dots + \Delta_n^{(\ell_n)} = T$. We then have the following result.

PROPOSITION 5. Let Assumptions 1, 2, and 3 hold. Let $\{\ell_n, \Delta_n^{(1)}, \dots, \Delta_n^{(\ell_n)}\}$ be defined as in (19) and (20) and put $\pi_n := \pi(\ell_n, \Delta_n^{(1)}, \dots, \Delta_n^{(\ell_n)})$, defined by Algorithm 3. Then, the sequence of policies $\{\pi_n\}$ is asymptotically optimal and satisfies

$$\sup_{\theta \in \Theta} \mathcal{R}_n^\pi(x, T; \theta) = O\left(\frac{(\log \log n)(\log n)^{1/2}}{n^{1/2}}\right). \quad (21)$$

Note that the regret of the sequence for policies introduced in Proposition 5 achieves the lower bound spelled out in Proposition 4 (up to logarithmic terms). In that sense, these policies cannot be improved upon. On the other hand, Algorithm 3 is restricted to the case where only one parameter is unknown, and exploits the fact that in this setting it is possible to learn the single unknown parameter by conducting price experiments in the neighborhood of the near-optimal price p^D . Designing policies that achieve the lower bound in Proposition 4 in the multiparameter case remains an open question.

6. Numerical Results and Qualitative Insights


6.1. Performance of the (Non)Parametric Policies and the “Price” of Uncertainty

We examine the performance of three policies developed in the previous sections: (i) the nonparametric policy defined in Algorithm 1; (ii) the parametric policy defined by means of Algorithm 2 and designed for a finite number of unknown parameters; and (iii) the parametric policy defined in Algorithm 3 designed for cases with a single unknown parameter. (The tuning parameters are taken as in Propositions 1, 3, and 5, respectively.)

The performance of these policies are measured by the magnitude of the regret. Note that $\mathcal{R}_n^\pi(x, T; \lambda) \approx C/n^\gamma$ implies that $\log \mathcal{R}_n^\pi(x, T; \lambda)$ should be approximately linear in $\log n$ with slope $-\gamma$. In Figure 1, we depict $\mathcal{R}_n^\pi(x, T; \lambda)$ as a function of n in a log-log plot for large values of n , and compute the best (least squares) linear fit for these values. The results depicted are based on running 10^3 independent simulation replications from which the performance indicators were derived by averaging. The standard error for $\mathcal{R}_n^\pi(x, T; \lambda)$ was below 0.05% in all cases.

Figure 1(a) summarizes results for an underlying exponential demand model $\lambda(p) = \theta \exp(-0.5p)$, where $\theta = 10 \exp(1)$, and Figure 1(b) presents results for an underlying linear demand model $\lambda(p) = 30 - \theta p$, where $\theta = 3$. The nonparametric algorithm does not make any assumptions with regard to the structure of the demand function (beyond those spelled out in §4), whereas the parametric algorithms are assumed to know the parametric structure, but the true value of θ is not revealed to them. In both cases, the initial normalized inventory level was $x = 20$, the selling horizon was $T = 1$, and the feasible price set was $[p, \bar{p}] = [0.1, 10]$.

Figure 1. Performance of pricing policies as a function of the market size (n).



Notes. Stars show the performance of the nonparametric policy defined in Algorithm 1; dots represent the performance of the parametric policy defined in Algorithm 2; and squares depict the performance of the parametric policy defined in Algorithm 3. The dashed lines represent the best linear fit to each set of points; in panel (a), the demand function is exponential and in panel (b) linear.

Discussion. The slopes of the best linear fit in Figures 1(a) and 1(b) are very close to $\gamma = -1/4$, $\gamma = -1/3$, and $\gamma = -1/2$, predicted by the upper bounds in (13), (17), and (21), respectively. These results provide a “picture proof” of the Propositions 1, 3, and 5. As is evident, the less structure is assumed a priori, the higher the profit loss relative to the full-information benchmark: informally speaking, this is the price paid due to increasing uncertainty with regard to the demand model.

An additional fundamental difference between the various policies concerns the degree to which the price domain is explored. The nonparametric policy (Algorithm 1) essentially needs to explore the entire price domain; the general parametric policy (Algorithm 2) needs to test a number of prices equal to the number of unknown parameters; and the parametric policy designed for the case of a single unknown parameter (Algorithm 3) explores only prices in the vicinity of the price p^D . In other words, as uncertainty decreases, the price exploration region shrinks as well.

6.2. The “Price” of Hedging Misspecification Risk

We have seen in the previous section that more refined information regarding the demand model yields higher revenues. Thus, it becomes tempting to postulate parametric structure, which in turn may be incorrect relative to the true underlying demand model. This misspecification risk can be eliminated via nonparametric pricing policies, but at

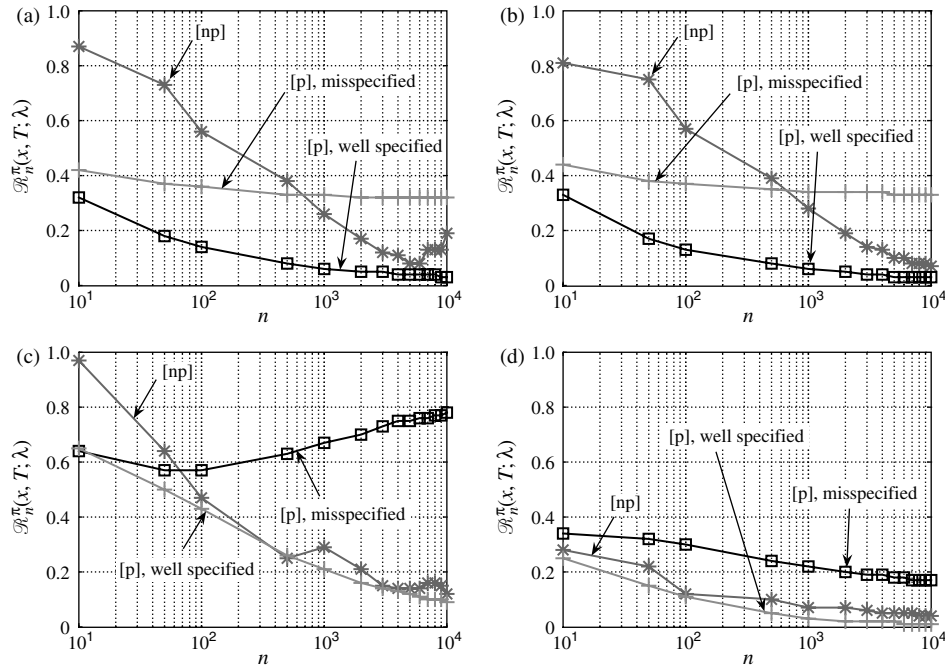
the price of settling for more modest performance revenue-wise. We now provide an illustration of this trade-off.

We fix the time horizon to be $T = 1$ and the set of feasible prices to be $[p, \bar{p}] = [0.1, 10]$. Figure 2 depicts the regret $\mathcal{R}_n^\pi(x, T; \lambda)$ for two underlying demand models, various values of n , and three policies. The first policy assumes an exponential parametric structure for the demand function, $\lambda(p; \theta) = \theta_1 \exp(-\theta_2 p)$ with $\theta = (\theta_1, \theta_2)$, and applies Algorithm 2 with $\tau_n = n^{-1/3}$. The second policy assumes a linear parametric structure $\lambda(p; \theta) = (\theta_1 - \theta_2 p)^+$ and applies Algorithm 2 with $\tau_n = n^{-1/3}$. Finally, the third policy is the one given by the nonparametric method described in Algorithm 1 (with tuning parameters as in Proposition 1). The point of this comparison is to illustrate that the parametric algorithm outperforms its nonparametric counterpart when the assumed parametric model is consistent with the true demand function; otherwise, the parametric algorithm leads to a regret that does not converge to zero due to a model misspecification error.

In Figures 2(a) and 2(b), the underlying demand model is given by $\lambda(p) = a \exp(-\alpha p)$ with $a = 10 \exp(1)$ and $\alpha = 1$, and hence the policy that assumes the exponential structure (squares) is *well specified*, whereas the policy that assumes a linear demand function (crosses) suffers from *model misspecification*.

In Figures 2(c) and 2(d), the underlying demand model is given by $\lambda(p) = (a - \alpha p)^+$ with $a = 30$ and $\alpha = 3$. Now the policy that assumes the exponential structure (squares) corresponds to a *misspecified* case, whereas the policy that

Figure 2. Performance of the parametric and nonparametric policies as a function of the market size (n).



Notes. The underlying demand model is exponential in Figures (a) and (b) and linear in Figures (c) and (d). Arrows indicate if the approach used is nonparametric [np] or parametric [p], and whether the model is well specified or misspecified in the case of the parametric algorithm. Crosses show the performance of the parametric policy (Algorithm 2) when a linear model is assumed while squares depict its performance when an exponential model is assumed; stars represent the performance of the nonparametric policy (Algorithm 1).

assumes a linear demand function (crosses) corresponds to a *well-specified* case. In Figures 2(a) and 2(c), the normalized inventory is $x = 8$, and in Figures 2(b) and 2(d), it is taken to be $x = 20$. The results depicted in the figures are based on running 10^3 independent simulation replications from which the performance indicators were derived by averaging. The standard error was below 0.9% in all cases.

Discussion. We focus on Figures 2(a) and 2(b) because similar remarks apply to Figures 2(c) and 2(d). First, note that the parametric algorithm is able to achieve close to 90% of the full-information revenues when the parametric model is well specified, when the market size is $n = 100$ or more. In addition, the convergence of the regret for Algorithm 2 is faster for small n under the well-specified parametric assumption (squares) in comparison to the nonparametric algorithm (stars). (We note that the nonmonotonic performance of the nonparametric algorithm (stars) stems from the fact that the price grid used by this algorithm changes with n , and consequently the minimal distance of the sought fixed price to any price on the grid need not be monotonic with respect to n .)

In contrast, if the model is misspecified, the regret $\mathcal{R}_n^\pi(x, T; \lambda)$ converges to a strictly positive value (crosses), as expected, and hence the performance of the parametric algorithm fails to asymptotically achieve the full-information revenues. The takeaway message here is that a nonparametric approach eliminates the risk stemming from

model misspecification, but at a price of extracting lower revenues than its parametric counterpart. The regret bounds derived in earlier sections precisely spell out the magnitude of this “price” of misspecification.

6.3. Numerical Analysis of the Bounds in Propositions 1 and 3

The results of Propositions 1 and 3 derive upper bounds on the regret that involve constants whose values have not been specified. Ideally, one would like to characterize the smallest-possible constants for which the upper bounds hold to determine the nonasymptotic behavior of the policies. Whereas an exact analysis along these lines is essentially intractable, it is possible to investigate this point numerically and assess the typical magnitude of these constants.

Fix $T = 1$ and $[p, \bar{p}] = [5, 10]$, and consider two families of demand functions: The first corresponds to an exponential class $\mathcal{L}_1 = \{a \exp\{-ap\} : a \in [\underline{a}, \bar{a}], \alpha \in [\underline{\alpha}, \bar{\alpha}]\}$, where $[\underline{a}, \bar{a}] = [5, 10]$ and $[\underline{\alpha}, \bar{\alpha}] = [0.1, 0.2]$; the second corresponds to a linear class $\mathcal{L}_2 = \{(b - \beta p)^+ : b \in [\underline{b}, \bar{b}], \beta \in [\underline{\beta}, \bar{\beta}]\}$, where $[\underline{b}, \bar{b}] = [10, 20]$ and $[\underline{\beta}, \bar{\beta}] = [0.2, 1]$. Note that all of the functions in the union of \mathcal{L}_1 and \mathcal{L}_2 are bounded by $M = \max\{\bar{a} \exp\{-\underline{\alpha} \bar{p}\}, \bar{b} - \underline{\beta} \bar{p}\} = 19$, are \bar{K} -Lipschitz with $\bar{K} = \max\{\bar{a} \bar{\alpha} \exp\{-\underline{\alpha} \bar{p}\}, \bar{\beta}\} = 1.21$, and have a first derivative bounded below by $\underline{K} = \min\{\underline{a} \alpha \exp\{-\bar{\alpha} \bar{p}\}, \underline{\beta}\} = 0.067$. Similarly, the maximum

Table 1. Estimates of the magnitude of the constants characterizing the performance of the algorithms: π_{np} corresponds to the nonparametric policy, and π_p to the parametric one.

	Demand functions					
	\mathcal{L}_1			\mathcal{L}_2		
	C_{\max}			C_{\max}		
Market size (n)	10^2	10^3	10^4	10^2	10^3	10^4
$x = 5$						
π_{np}	0.92	1.04	1.07	1.09	1.25	1.31
π_p	0.79	0.79	0.79	1.11	1.11	1.11
$x = 10$						
π_{np}	0.66	0.75	0.77	0.74	0.85	0.87
π_p	0.61	0.53	0.39	0.58	0.60	0.54

revenue rate is bounded below by $m = \min\{\underline{p}a \exp\{-\bar{\alpha}\bar{p}\}, \underline{p}(\underline{b} - \bar{\beta}\underline{p})\} = 3.38$.

To estimate the constants characterizing the revenue losses of the parametric and nonparametric policies, we proceed as follows. We draw 100 demand functions from each class of demand models above, using a uniform distribution over each interval defining the space of parameters. For each function, we estimate the regret associated with the proposed nonparametric (Algorithm 1) and parametric (Algorithm 2) policies for market sizes of $n = 100$, $n = 10^3$, and $n = 10^4$, and for two initial (normalized) inventory levels ($x = 5$, $x = 10$). This step is executed as in the previous subsections. For each draw and each market size, we compute the value of $\mathcal{R}_n^\pi(x, T; \lambda)n^\gamma$, where $\gamma = 1/4$ for the nonparametric policy, and $\gamma = 1/3$ for the parametric one. Taking the maximum of these values over all 100 draws, we arrive at an estimate, C_{\max} , for each market size, which corresponds to the worst-case constant observed for all models. This value serves as a proxy for the magnitude of the constants characterizing the performance of the proposed algorithms. Results are summarized in Table 1.

The reasonable magnitude of the constants (all fall in the range of 0.5–1.3) suggests that our policies should perform well in practical settings. In particular, based on these results, we expect the nonparametric policy to achieve at most a regret of 35% for $n = 100$, 23% for $n = 10^3$, and 14% for $n = 10^4$, whereas the parametric policy should achieve a regret of 24% for $n = 100$, 12% for $n = 10^3$, and 6% for $n = 10^4$.

7. Selected Proofs

Preliminaries and Notation. For any real number x , x^+ will denote $\max\{x, 0\}$. C_1, C_2, \dots will be used to denote positive constants that are independent of a given demand function, but may depend on the parameters of the class \mathcal{L} of admissible demand functions, and on x and T . A sequence $\{a_n\}$ of real numbers is said to increase to

infinity at a *polynomial rate* if there exists a constant $\beta > 0$ such that $\liminf_{n \rightarrow \infty} a_n/n^\beta > 0$. To lighten notation, we will occasionally omit the arguments of $J_n^\pi(x, T; \lambda)$ and $J_n^D(x, T | \lambda)$.

PROOF OF PROPOSITION 1. The proof is organized around three main steps that develop and analyze a lower bound on the expected revenues achieved by the proposed sequence of policies $\pi_n := \pi(\tau_n, \kappa_n)$, $n = 1, 2, \dots$ defined by means of Algorithm 1.

Fix $\lambda \in \mathcal{L}$, $\eta = 2$. Let τ_n be such that $\tau_n \rightarrow 0$ and $n\tau_n \rightarrow \infty$ at a polynomial rate as $n \rightarrow \infty$. Let κ_n be a sequence of integers such that $\kappa_n \rightarrow \infty$ and $n\Delta_n := n\tau_n/\kappa_n \rightarrow \infty$ at a polynomial rate. We divide the interval $[\underline{p}, \bar{p}]$ into κ_n equal-length intervals, and we let $P_n = \{p_i, i = 1, \dots, \kappa_n\}$ be the left endpoints of these intervals. Now partition $[0, \tau_n]$ into κ_n intervals of length Δ_n and apply the price p_i on the i th interval. Define

$$\hat{\lambda}(p_i) = \frac{N(\sum_{j=1}^i n\lambda(p_j)\Delta_n) - N(\sum_{j=1}^{i-1} n\lambda(p_j)\Delta_n)}{n\Delta_n}, \quad i = 1, \dots, \kappa_n, \quad (22)$$

where $N(\cdot)$ is a unit-rate Poisson process. Thus, $\hat{\lambda}(p_i)$ denotes the number of product requests over successive intervals of length Δ_n , normalized by $n\Delta_n$. Put

$$u_n := (\log n)^{1/2} \max\{1/\kappa_n, 1/(n\Delta_n)^{1/2}\}. \quad (23)$$

This will be used to quantify deviations of various quantities associated with the proposed policy from their full-information counterparts.

Step 1. We first develop an expression for a lower bound on the expected revenues achieved by the proposed policy π_n . Let $X_n^{(L)} = \sum_{i=1}^{\kappa_n} \lambda(p_i)n\Delta_n$, $X_n^{(P)} = \lambda(\hat{p})n(T - \tau_n)$ and put $Y_n = N(X_n^{(L)} + X_n^{(P)})$, $Y_n^{(L)} = N(X_n^{(L)})$, and $Y_n^{(P)} = Y_n - Y_n^{(L)}$. $Y_n^{(L)}$ represents the maximum number of requests during the learning phase if the system would not run out of resources, $Y_n^{(P)}$ the maximum number of requests during the pricing phase, and Y_n the maximum total number of requests throughout the sales horizon. Now note that one can lower bound the revenues achieved by π_n by those accumulated during the pricing phase, and during the latter, the maximum number of units that can be sold is exactly $\min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\}$. We deduce that

$$J_n^\pi \geq \mathbb{E}[\hat{p} \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\}]. \quad (24)$$

Next, we analyze the lower bound above by getting a handle on \hat{p} , the estimate of p^D and the quantities $Y_n^{(P)}$ and $Y_n^{(L)}$.

Step 2. In what follows, we separate two cases: $\lambda(\bar{p}) \leq x/T$ and $\lambda(\bar{p}) > x/T$. The latter case is one where the decision maker would on average run out of inventory during the sales horizon, and has to be analyzed separately.

Case 1. Suppose first that $\lambda(\bar{p}) \leq x/T$. Note that $\min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\} = Y_n^{(P)} - (Y_n - nx)^+$ and that $\mathbb{E}[\hat{p}Y_n^{(P)}] = \mathbb{E}[r(\lambda(\hat{p}))n(T - \tau_n)]$. Hence, using (24) and the fact that $\hat{p} \leq \bar{p}$, we have

$$\begin{aligned} J_n^\pi &\geq \mathbb{E}[\hat{p}Y_n^{(P)}] - \bar{p}\mathbb{E}[(Y_n - nx)^+] \\ &= \mathbb{E}[r(\lambda(\hat{p}))n(T - \tau_n) - \bar{p}\mathbb{E}[(Y_n - nx)^+]]. \end{aligned} \quad (25)$$

Lemma 4 in Online Appendix C establishes that $\mathbb{E}[r(\lambda(\hat{p}))]$ is close to $r(\lambda(p^D))$, namely, $\mathbb{E}[r(\lambda(\hat{p}))] \geq r(\lambda(p^D)) - C_1u_n - C_2/n^\eta$ for positive constants C_1, C_2 . This is done by controlling the deviations of a Poisson process from its mean. With respect to the second term on the right-hand side of (25), Lemma 5 in Online Appendix C establishes that $\mathbb{E}[(Y_n - nx)^+] \leq C_3nu_n$. The key issue here is to control for stochastic fluctuations of customer requests at the estimated price \hat{p} . We thus conclude that

$$\begin{aligned} J_n^\pi &\geq \left[r(\lambda(p^D)) - C_1u_n - \frac{C_2}{n^\eta} \right] n(T - \tau_n) - \bar{p}C_3nu_n \\ &\geq nr(\lambda(p^D))T - C_4n(u_n + \tau_n), \end{aligned} \quad (26)$$

where C_4 is a suitably chosen constant. Lemma 3 in Online Appendix C establishes a lower bound on $J_n^D, J_n^D \geq nm^D = nm \min\{T, x/M\}$, where m and M were defined in Assumption 1. The latter, in conjunction with (26), implies that

$$\frac{J_n^\pi}{J_n^D} = \frac{J_n^\pi}{nJ^D} \geq 1 - \frac{C_4}{m^D}(u_n + \tau_n). \quad (27)$$

Case 2. Now suppose that $\lambda(\bar{p}) > x/T$. In this case, note that $p^c = p^D = \bar{p}$. Recalling the lower-bound expression in (24), we seek to get a handle on $\min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\}$. Lemma 6 in Online Appendix C establishes that this quantity will be close to the total units in inventory, nx , with very high probability. This is done by bounding the probability of $\mathcal{A} := \{\omega: \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\} \geq nx - C_5u_n, |\hat{p} - p^D| \leq C_6u_n\}$ for appropriate positive constants C_5, C_6 . The revenues generated by π_n can be bounded below as follows:

$$\begin{aligned} J_n^\pi &\geq \mathbb{E}[\hat{p} \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\}] \\ &\stackrel{(a)}{\geq} \mathbb{E}[(p^D - C_6u_n) \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\} | \mathcal{A}] \mathbb{P}(\mathcal{A}) \\ &\stackrel{(b)}{\geq} (p^D - C_6u_n)(nx - C_9nu_n) \left(1 - \frac{C_7}{n^{\eta-1}}\right) \\ &\geq p^D nx - nC_8u_n, \end{aligned} \quad (28)$$

where both (a) and (b) follow from the definition of \mathcal{A} and Lemma 6, and where $C_7, C_8 > 0$ are suitably large. Now, recall that because $\lambda(\bar{p}) > x/T$, $p^D = \bar{p}$ and $J_n^D = nx\bar{p}$. Hence,

$$\frac{J_n^\pi}{J_n^D} = \frac{J_n^\pi}{nx\bar{p}} \geq 1 - \frac{C_8}{\bar{p}x}u_n. \quad (29)$$

Step 3. Let $C_9 = \max\{C_4/(m^D), C_8/(\bar{p}x)\}$. Combining (27) and (29), we have for all $\lambda \in \mathcal{L}$, $\mathcal{R}_n^\pi(x, T; \lambda) \leq C_9(u_n + \tau_n)$, and note that C_9 does not depend on the specific function $\lambda \in \mathcal{L}$. Note that the choice of tuning parameters that minimizes the order of the the upper bound on the regret is exactly $\tau_n \asymp n^{-1/4}$ and $\kappa_n \asymp n^{1/4}$. Plugging in, we get for some $C_{10} > 0$,

$$\sup_{\lambda \in \mathcal{L}} \mathcal{R}_n^\pi(x, T; \lambda) \leq \frac{C_{10}(\log n)^{1/2}}{n^{1/4}}.$$

The proof is complete. \square

PROOF OF PROPOSITION 2. Because $J_n^\pi(x, T; \lambda) \leq J_n^* \cdot (x, T | \lambda)$, the result will be established if we can prove that for some $C > 0$

$$\sup_{\lambda \in \mathcal{L}} \left(1 - \frac{J_n^*(x, T | \lambda)}{J_n^D(x, T | \lambda)}\right) > \frac{C}{n^{1/2}}. \quad (30)$$

Consider the demand function $\lambda(p) = (a - bp)^+$ with $b \in [\underline{K}, \bar{K}]$ and $a = \max\{2\bar{K}\bar{p}, \bar{K}\bar{p} + x/T\}$. Note in particular that $\lambda \in \mathcal{L}$ because $a \leq M$. In addition, for this demand function we have $p^u = p^c = \bar{p}$. By Proposition 1 in Gallego and van Ryzin (1994), any dynamic pricing policy will never price below p^u ; hence, the optimal dynamic pricing policy is to price at p^u until the minimum of T and the time when inventory is depleted. We deduce that

$$\begin{aligned} J_n^*(x, T | \lambda) &= p^u \mathbb{E}[N(n\lambda(p^u)T)] \\ &\quad - p^u \mathbb{E}[(N(n\lambda(p^u)T) - nx)^+] \\ &= nr(p^u)T - p^u \mathbb{E}[(N(nx) - nx)^+] \\ &= J_n^D(x, T | \lambda) - \bar{p}nx e^{-nx} \frac{(nx)^{nx}}{(nx)!}. \end{aligned}$$

Now, using, e.g., Sterling's approximation, one has that

$$\begin{aligned} n^{1/2}(1 - J_n^*/J_n^D) &= (n^{1/2}\bar{p}\lambda(\bar{p})T)^{-1}\bar{p}nx \\ &\quad \cdot e^{-nx}(nx)^{nx}((nx)!)^{-1} \rightarrow C_1 > 0 \end{aligned}$$

as $n \rightarrow \infty$ and (30) follows. This completes the proof. \square

PROOF OF PROPOSITION 3. In this section, we define for any $\theta \in \Theta$, $p^D(\theta) := \max\{p^u(\theta), p^c(\theta)\}$, where $p^u(\theta) = \arg \max\{p\lambda(p; \theta): p \in [\underline{p}, \bar{p}]\}$, $p^c(\theta) = \arg \min\{|\lambda(p; \theta) - x/T|: p \in [\underline{p}, \bar{p}]\}$. In addition, we define $\mathcal{D}_\lambda(\theta) = \{l: l = \lambda(p; \theta) \text{ for some } p \in [\underline{p}, \bar{p}]\}$. Let θ^* denote the true parameter value. The proof follows the structure of that of Proposition 1. The first step provides a lower bound on the revenues achieved by the proposed policy. The second step bounds the difference between the estimated parameter vector and its true value and analyzes the revenue rate at $\hat{p} = \max\{p^u(\hat{\theta}), p^c(\hat{\theta})\}$. The proof concludes with the resulting upper bound obtained for the regret.

Step 1. Here, we derive a lower bound on the expected revenues under the policy π_n . Let $X_n^{(L)} = \sum_{i=1}^k \lambda(p_i)n\Delta_n$, $X_n^{(P)} = \lambda(\hat{p})n(T - \tau_n)$, and put $Y_n = N(X_n^{(L)} + X_n^{(P)})$, $Y_n^{(L)} = N(X_n^{(L)})$, and $Y_n^{(P)} = Y_n - Y_n^{(L)}$. As in the proof of Proposition 1, we have

$$J_n^\pi \geq \mathbb{E}[\hat{p} \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\}] \quad (31)$$

Next, we analyze the lower bound above by getting a handle on \hat{p} , the estimate of p^D , and the quantities $Y_n^{(P)}$ and $Y_n^{(L)}$.

Step 2. Lemma 7 in Online Appendix C establishes that $r(\lambda(p^D(\theta^*); \theta^*); \theta^*) - r(\lambda(\hat{p}; \theta^*); \theta^*) \leq C_1 \|\hat{\theta} - \theta^*\|_\infty$, which implies that controlling $\hat{\theta}$ will provide a handle on \hat{p} . Lemma 8 in Online Appendix C establishes that $\mathbb{E}\|\hat{\theta} - \theta^*\|_\infty \leq C_2(n\tau_n)^{-1/2}$ for all $n \geq 1$, which in turn implies that

$$\begin{aligned} \mathbb{E}[r(\lambda(\hat{p}; \theta^*); \theta^*)] &= r(\lambda(p^D(\theta^*); \theta^*); \theta^*) + \mathbb{E}[r(\lambda(\hat{p}; \theta^*); \theta^*) \\ &\quad - r(\lambda(p^D(\theta^*); \theta^*); \theta^*)] \\ &\geq r(\lambda(p^D(\theta^*); \theta^*); \theta^*) - \frac{C_1 C_2}{(n\tau_n)^{1/2}}. \end{aligned}$$

Finalizing the analysis of the lower bound in (31) can be conducted in a similar manner as in the proof of Proposition 1 (Step 2) by letting in this case $u_n = (\log n)^{1/2}(n\tau_n)^{-1/2}$. Indeed, one can again focus on the two cases $\lambda(\bar{p}; \theta^*) \leq x/T$ and $\lambda(\bar{p}; \theta^*) > x/T$. In the former case, the key is on getting a handle of $\mathbb{E}[(Y_n - nx)^+]$, which can be shown to be bounded by $C_3 u_n$ for some $C_3 > 0$. In the latter, one can establish that the firm will sell almost all units with very high probability. In particular, one can show that $\mathcal{A} := \{\omega: \min\{Y_n^{(P)}, (nx - Y_n^{(L)})^+\} \geq nx - C_4 u_n, |\hat{p} - p^D| \leq C_5 u_n\}$ has very high probability for an appropriate choice of $C_4, C_5 > 0$.

Step 3. As in the proof of Proposition 1 (Step 3), one obtains that for some $C_6 > 0$,

$$\left(1 - \frac{J_n^\pi}{J_n^D}\right) \leq C_6(\tau_n + (\log n)^{1/2}(n\tau_n)^{-1/2}), \quad (32)$$

and the result follows by plugging in $\tau_n \asymp n^{-1/3}$. This completes the proof. \square

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Acknowledgments

This research was partially supported by NSF grant DMI-0447562.

References

- Araman, V. F., R. A. Caldentey. 2005. Dynamic pricing for non-perishable products with demand learning. Working paper, New York University.
- Auer, P., N. Cesa-Bianchi, Y. Freund, R. E. Schapire. 2002. The non-stochastic multiarmed bandit problem. *SIAM J. Comput.* **32** 48–77.
- Aviv, Y., A. Pazgal. 2005. Pricing of short life-cycle products through active learning. Working paper, Washington University, St. Louis.
- Ball, M., M. Queyranne. 2009. Toward robust revenue management: Competitive analysis of online booking. *Oper. Res.* **57**(4) 950–963.
- Bertsimas, D., G. Perakis. 2006. Dynamic pricing: A learning approach. S. Lawphongpanich, D. W. Hearn, M. J. Smith, eds. *Mathematical and Computational Models for Congestion Charging*. Springer US, 45–79.
- Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5** 203–229.
- Borovkov, A. 1998. *Mathematical Statistics*. Gordon and Breach Science Publishers.
- Carvalho, A. X., M. L. Puterman. 2005. Dynamic pricing and learning over short time horizons. Working paper, University of British Columbia, Vancouver, British Columbia, Canada.
- Cesa-Bianchi, N., G. Lugosi. 2006. *Prediction, Learning, and Games*. Cambridge University Press, Cambridge, UK.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* **49** 1287–1309.
- Eren, S., C. Maglaras. 2007. Pricing without market information. Working paper, Columbia University, New York.
- Farias, V. F., B. Van Roy. 2006. Dynamic pricing with a prior on market response. Working paper, Stanford University, Stanford, CA.
- Foster, D. P., R. Vohra. 1999. Regret in the on-line decision problem. *Games Econom. Behav.* **29** 7–35.
- Gallego, G., G. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **50** 999–1020.
- Hannan, J. 1957. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, Vol. 3. Princeton University Press, Princeton, NJ, 97–139.
- Huh, W. T., P. Rusmevichientong. 2008. Online sequential optimization with biased gradients: Theory and applications to censored demand. Working paper, Columbia University, New York.
- Kleinberg, R., F. T. Leighton. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. *Proc. 44th Ann. IEEE Sympos. Foundations Comput. Sci.* IEEE, Piscataway, NJ, 594–605.
- Lai, T. L., H. Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* **6** 4–22.
- Lim, A., J. Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* **55**(2) 198–214.
- Lobo, M. S., S. Boyd. 2003. Pricing and learning with uncertain demand. Working paper, Duke University, Durham, NC.
- Perakis, G., G. Roels. 2009. Distribution-free revenue management. *Manufacturing Service Oper. Management*. Forthcoming.
- Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **58** 527–535.
- Rusmevichientong, P., B. van Roy, P. W. Glynn. 2006. A non-parametric approach to multi-product pricing. *Oper. Res.* **54** 82–98.
- Scarf, H. 1959. Bayes solutions of the statistical inventory problem. *Ann. Math. Statist.* **30** 490–508.
- Talluri, K. T., G. J. van Ryzin. 2005. *Theory and Practice of Revenue Management*. Springer-Verlag.
- van Ryzin, G., J. McGill. 2000. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Sci.* **46** 760–775.